EXHIBIT 1

UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

VIRTAMOVE, CORP., Plaintiff, v. HEWLETT PACKARD ENTERPRISE COMPANY, Defendant.	§ § § § § § § § § § §	Case No. 2:24-cv-00093-JRG (Lead Case) **JURY TRIAL DEMANDED**
VIRTAMOVE, CORP., Plaintiff, v. INTERNATIONAL BUSINESS MACHINES CORP., Defendant.	§ § § § § § § § § § §	Case No. 2:24-cv-00064-JRG (Member Case) **JURY TRIAL DEMANDED**

DECLARATION OF DR. ANGELOS STAVROU

I. INTRODUCTION

1. My name is Dr. Angelos Stavrou. I have been retained to provide technical assistance in the above-captioned action on behalf of Defendants International Business Machines, Corp. ("IBM") and Hewlett Packard Enterprise Company ("HPE"). I have been asked to offer my opinions relating to the meaning of the following terms used in the claims of U.S. Patent No. 7,784,058 (the "'058 Patent") to a person of ordinary skill in the art, including whether, when read in light of the specifications and the prosecution history, they inform those skilled in the art of the scope of the invention with reasonable certainty:

- critical system elements / operating system critical system elements / shared library critical system elements ['058 Patent claims 1, 2, 3, 4, and 18];
- functional replicas ['058 Patent claim 1].

2. I based the opinions provided in this Declaration on my years of experience in the field, education, and review of the materials identified in this Declaration and in Appendix A to this Declaration.

3. I reserve the right to amend or supplement my opinions in light of further documents, depositions, or discovery disclosures. I further reserve the right to rely upon any additional information or materials that may be provided to me or that are relied upon by any of Plaintiff VirtaMove, Corp.'s ("Plaintiff") experts or witnesses, if called to testify or to give additional opinions regarding this matter. I specifically reserve the right to address, including through reference to new evidence and/or opinions, the arguments and opinions raised by Plaintiff or its experts once they are disclosed, such as in any expert report or declaration.

II. QUALIFICATIONS AND EXPERIENCE

4. My qualifications, including my education and work experience, are provided in my Curriculum Vitae, which is attached as Appendix B and summarized as follows.

5. I graduated from the University of Patras in 1997 with a B.Sc. (Honors) degree in Physics, and with a certificate in Electrical Engineering in 1999. I received a Master of Science in Electrical Engineering and a Master of Philosophy in Computer Science from Columbia University, Fu Foundation School of Engineering and Applied Science in 2002 and 2007 respectively. I received my Ph.D. with distinction in Computer Science from Columbia University in 2007. I am currently a professor at Virginia Polytechnic Institute and State University ("Virginia Tech") in the Bradley Department of Electrical and Computer Engineering, and I have been a professor of computer science and engineering since 2007, including at George Mason University in the Computer Science Department.

6. I have authored or co-authored over 120 peer-reviewed conference and journal articles on various topics relating to electrical and computer engineering, including operating systems, virtualization technologies, networking, systems security, and data and resource management. *See, e.g.*:

- Xing Gao, Jidong Xiao, Haining Wang & Angelos Stavrou, *Understanding the Security Implication of Aborting Virtual Machine Live Migration*, IEEE TRANSACTIONS ON CLOUD COMPUTING (Apr–June 2022);
- Fengwei Zhang, Kevin Leach, Angelos Stavrou & Haining Wang, *Towards Transparent Debugging*, IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING (Mar.–Apr. 2018);
- Fengwei Zhang, Jiang Wang, Kun Sun & Angelos Stavrou, *Hypercheck: A Hardware-Assisted Integrity Monitor*, IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING (July–Aug. 2014);
- Meixing Le. Angelos Stavrou & Brent ByungHoon Kang, *DoubleGuard: Detecting Intrusions in Multitier Web Applications*, IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING (July–Aug. 2012);
- Tolga O. Atalay, Sudip Maitra, Dragoslav Stojadinovic, Angelos Stavrou & Haining Wang, *Securing 5G OpenRAN with a Scalable Authorization Framework for xApps* (IEEE Conference on Computer Communications, 2023);

- Tolga O. Atalay, Dragoslav Stojadinovic, Alireza Famili, Angelos Stavrou & Haining Wang, *Network-Slice-as-a-Service Deployment Cost Assessment in an End-to-End 5G Testbed* (IEEE Global Communications Conference, 2022);

- Tolga O. Atalay, Dragoslav Stojadinovic, Angelos Stavrou & Haining Wang, *Scaling Network Slices with a 5G Testbed: A Resource Consumption Study* (IEEE Wireless Communications and Networking Conference, 2022);

- Joseph Connelly, Taylor Roberts, Xing Gao, Jidong Xiao, Haining Wang & Angelos Stavrou, *CloudSkulk: A Nested Virtual Machine Based Rootkit and Its Detection* (51st Annual IEEE/IFIP International Conference on Dependable Systems, 2021);

- Jiang Wang, Kun Sun & Angelos Stavrou, *Hardware-Assisted Application Integrity Monitor* (Hawaii International Conference on System Sciences, 2012);

- Jiang Wang, Sameer Niphadkar, Angelos Stavrou & Anup K. Ghosh, *A Virtualization Architecture for In-Depth Kernel Isolation* (Hawaii International Conference on System Sciences, 2010).

7. I have also authored several book chapters in the areas of networking, systems security, and application isolation technologies. *See, e.g.*:

- Angelos Stavrou, *Overlay-Based DoS Defenses*, *in* ENCYCLOPEDIA OF CRYPTOGRAPHY AND SECURITY (Henk C.A. van Tilborg & Sushil Jajodia eds., 2d ed. 2010);

- Angelos Stavrou, *TCP Modulation Attacks*, *in* ENCYCLOPEDIA OF CRYPTOGRAPHY AND SECURITY (Henk C.A. van Tilborg & Sushil Jajodia eds., 2d ed. 2010).

8. I have also hosted and participated in several workshops on virtualization technologies, networking, and systems security. *See, e.g.*:

- Ataollah Fatahi Baarzi, George Kesidis, Dan Fleck & Angelos Stavrou, *Microservices made attack-resilient using unsupervised service fissioning* 31 (13th European Workshop on Systems Security, 2020);

- Riyadh Mahmood, Naeem Esfahani, Thabet Kacem, Nariman Mirzaei, Sam Malek & Angelos Stavrou, *A Whitebox Approach for Automated Security Testing of Android Applications on the Cloud* 1 (7th International Workshop on Automation of Software Test, 2012);

- Angelos D. Keromytis, Roxana Geambasu, Simha Sethumadhavan, Salvatore J. Stolfo, Junfeng Yang,Azzedine Benameur, Marc Dacier, Matthew Elder, Darrell Kienzle & Angelos Stavrou, *The MEERKATS Cloud Security Architecture* 446 (3rd International Workshop on Security and Privacy in Cloud Computing, 2012);

- Angelos D. Keromytis, Salvatore J. Stolfo, Junfeng Yang, Angelos Stavrou, Anup Ghosh, Dawson Engler, Marc Dacier, Matthew Elder & Darrell Kienzle, *The MINESTRONE Architecture Combining Static and Dynamic Analysis Techniques for Software Security* 53 (1st SysSec Workshop, 2011);

- Charalampos Andrianakis, Paul Seymer & Angelos Stavrou, *Scalable Web Object Inspection and Malfease Collection* 1 (5th USENIX Workshop on Hot Topics in Security, 2010);

- Quan Jia, Zhaohui Wang & Angelos Stavrou, *The Heisenberg Measuring Uncertainty in Lightweight Virtualization Testbeds* 1 (2nd Workshop on Cyber Security Experimentation and Test, 2009);

- Yih Huang, Angelos Stavrou, Anup K. Ghosh & Sushil Jajodia, *Efficiently Tracking Application Interactions using Lightweight Virtualization* 19 (1st Workshop on Virtualization Security, 2008);

- Michael E. Locasto, Angelos Stavrou & Angelos D. Keromytis, *Dark Application Communities* 11 (15th New Security Paradigms Workshop, 2006).

9. I have been involved in a variety of research efforts in the area of electrical and computer engineering, including operating systems, networking, systems security, data and resource management, and virtualization technologies, since 2007. My research has been supported by awards and grants from the National Science Foundation, Defense Advanced Research Projects Agency (DARPA), the National Institute of Standards and Technology (NIST), the Army Research Office (ARO), the Department of Homeland Security, and other institutions. *See, e.g.*:

- Principal Investigator (Kryptowire and Virginia Tech) for DARPA Cyber Agents for Security Testing and Learning Environments (CASTLE), 06/01/2023 – ongoing, Multi-Agent Training Exerciser (MATrEx) along with SRC and RTX, total budget $16.5M;

- Principal Investigator (George Mason University) for DARPA Mission-oriented Resilient Clouds (MRC) project, $750,363 09/2011–01/2016, *MEERKATS: Maintaining Enterprise Resiliency via Kaleidoscopic Adaptation & Transformation of Software Services*, (Part of team that includes Columbia University and Symantec Corp. total budget: $6,619,270) (with Fei Li);

- Principal Investigator, NSF project, $239,884 09/2009–08/2011, *TC: Small: Collaborative Research: Scalable Malware Analysis Using Lightweight Virtualization* (with Fabian Monrose).

10. I am also a named inventor on several issued patents relating to operating systems, networking, systems security, data and resource management, and virtualization technologies. *See, e.g.*:

- U.S. Patent No. 10,127,137 (*Methods and systems for increasing debugging transparency*);
- U.S. Patent No. 9,602,524 (*Methods and Apparatus for Application Isolation*);
- U.S. Patent No. 9,098,698 (*Methods and Apparatus for Application Isolation*);
- U.S. Patent No. 9,270,697 (*Hardware-assisted Integrity Monitor*);
- U.S. Patent No. 8,819,225 (*Hardware-assisted Integrity Monitor*);
- U.S. Patent No. 8,549,646 (*Methods, Media and Systems for Responding to a Denial of Service Attack*).

11. I have been a member of the Institute of Electrical and Electronics Engineers ("IEEE"), since 2007 and am now a Senior Member. I have served in various editorial bodies, including as an editor for the IEEE Security and Privacy Magazine, IEEE Transactions on Computers, IEEE Internet Computing, and IEEE Transactions on Reliability. I have also served on the IEEE Rebooting Computing Committee since 2013. In 2024, I received the IEEE Reliability Society Lifetime Achievement Award. Additionally, I have served as one of the chairs for programs such as the Association for Computing Machinery ("ACM") Conference on Computer and Communications Security (CCS) and the 1st Workshop on Virtual Machine Security (VMSec).

12. In sum, I have almost two decades of experience in research and development in the areas of electrical and computer engineering as a professor, researcher, and consultant.

13. I am being compensated at my usual hourly rate of $450. My compensation does not depend on the outcome of these proceedings. I have no financial interest, beneficial or otherwise, in any of the parties.

III. LEGAL PRINCIPLES

14. I am not a lawyer and offer no opinions on the law. For the purposes of performing my analyses, I have been instructed by counsel of certain legal principles with respect to patent law and claim construction, and I have applied those principles in reaching the opinions set forth herein.

A. Person of Ordinary Skill in the Art

15. I understand that the claims of a patent define the purported invention, and the purpose of claim construction is to understand how a person of ordinary skill in the art ("POSITA") would have understood disputed claim terms at the time of the purported invention. I understand factors that may be considered in determining the level of ordinary skill in the art include: (a) the types of problems encountered in the field, (b) prior art solutions to those problems, (c) how sophisticated the technology at issue is and how fast inventions occur in the field, and (d) the educational level of active workers in the field.

16. As to the '058 Patent, a POSITA would have a bachelor's degree in computer science or a similar degree, and at least two years of experience in the field of computer software systems. Additional education might compensate for less experience, and vice-versa.

17. My opinions in this Declaration are from the perspective of a POSITA. Given my education and experience, I was a POSITA as of the earliest priority date of the '058 Patent.

B. Claim Construction

18. I understand that the most important evidence to consider in construing the claims is the "intrinsic" record. This includes the claim language, the patent specification, and the prosecution history. I understand that words or terms should be given their ordinary and accepted meaning to a POSITA unless it appears that the inventors were using them to mean something else. In making this determination, however, of paramount importance are the claims, the patent

specification, and the prosecution history. Additionally, the specification and prosecution history should be consulted to confirm whether the patentee has acted as its own lexicographer (*i.e.*, provided a special meaning to any disputed terms), or intentionally disclaimed, disavowed, or surrendered any claim scope.

19. I further understand that reference materials that were publicly available at the time that the patent application was filed, such as dictionaries, treatises, or other technical references, may provide context and background for understanding how a POSITA would have considered the terms used in the claims. However, I understand that such references, as well as testimony (including this Declaration), are generally known as "extrinsic evidence," and are accorded less weight than evidence found within the patent and prosecution history.

20. I understand that POSITAs are deemed to read the words used in the patent with an understanding of their meaning in the field, and to have and apply knowledge of any special meaning and usage associated with the term within the field. The words used by the inventors to describe the invention must be understood and interpreted as they would be understood and interpreted by a POSITA.

21. I understand that the claims of a patent define the scope of the rights conferred by the patent. I further understand that the claims of a patent must satisfy a definiteness requirement, i.e., the claims must particularly point out and distinctly claim the subject matter which the patentee regards as his invention. I also understand that a patent is invalid for indefiniteness if its claims, read in light of the patent specification and the file history, fail to inform, with reasonable certainty, those skilled in the art about the scope of the claimed subject matter. It was further explained to me that absolute or mathematical precision in claim language is not required. However, it is not enough that some meaning can be ascribed to a patent's claims, including where, for example, the

patent specification purports to provide a definition for a claim term. Instead, the claims, when read in light of the patent specification and prosecution history, must provide objective boundaries for a POSITA.

22. I have also been informed that terms of degree used in a patent are indefinite unless the patent provides objective boundaries for those of ordinary skill in the art when read in light of the intrinsic evidence.

IV. THE '058 PATENT

23. The '058 Patent is titled "Computing System Having User Mode Critical System Elements as Shared Libraries," and it purports to recite "[a] computing system and architecture . . . that affects and extends services exported through application libraries." '058 Patent at Abstract, 1:15-18.

24. The '058 Patent explains that "[c]omputer systems are designed in such a way that software application programs share common resources," and "it is traditionally the task of an operating system to provide mechanisms to safely and effectively control access to shared resources." '058 Patent at 1:21-24. The '058 Patent further explains that "the centralized control" of shared resources by the operating system "creates a limitation caused by conflicts for shared resources." '058 Patent at 1:26-28.

25. The '058 Patent purports to address the problem where conflicts are created by multiple software applications accessing the same kernel level services on a single computing platform by making such services available in user-mode shared libraries, which the software applications can access thereafter. '058 Patent at 1:21-54; 6:62-8:44. Specifically, the '058 Patent describes replicating CSEs "through the use of shared libraries" that "provide[] a means of attaching or linking a CSE service to an application having access to the shared library." '058

Patent at 5:29-34. These replicated CSEs are called "shared library critical system elements" ("SLCSEs") and are located in user space, outside of the operating system. '058 Patent at 2:10, 3:19-24, 5:21-28.

26. An exemplary embodiment is presented in Figure 3 and Figure 4:

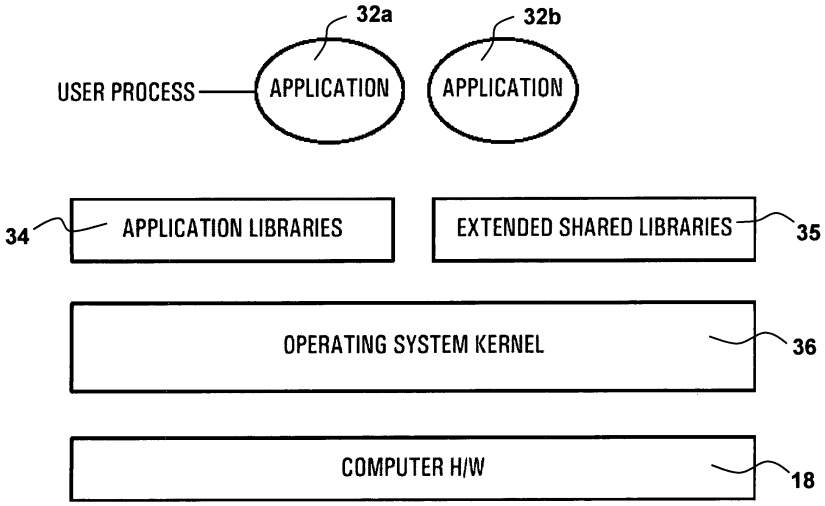

FIG. 3

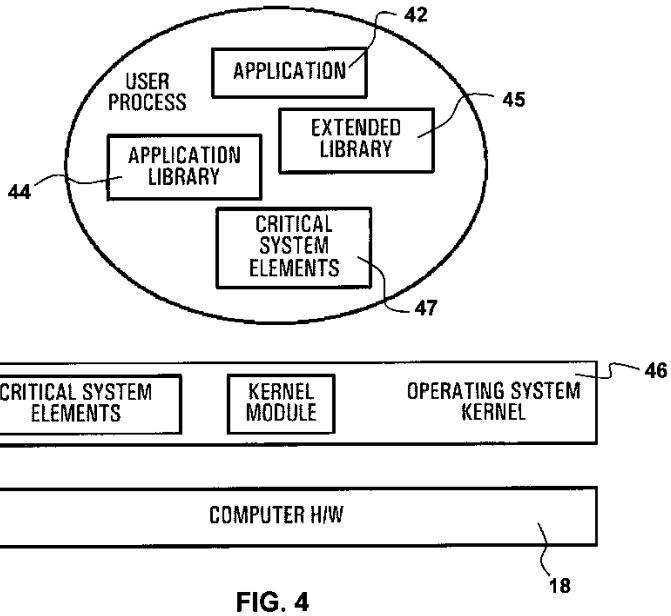

FIG. 4

27. As shown in Figure 3, software applications utilize their own application libraries, which "are augmented by an extension, which contains critical system elements, that reside in extended shared libraries." '058 Patent at 8:11-15. The '058 Patent explains that "some system elements that are critical to the operation of a software application are replicated from kernel mode, into user mode in the same context as that of the application" and that "[t]hese system elements are contained in a shared library." '058 Patent at 9:15-19. The '058 Patent explains "[t]he process of starting the application includes creating a linkage to all of the required shared libraries that the application will need," such that the "CSE is loaded and linked to the application." '058 Patent at 9:22-26. The '058 Patent also explains that "[e]ach application that will use a CSE will link to a library containing the CSE independent of any other application" in order to "provide an application with a unique instance of a CSE." '058 Patent at 3:26-30.

28. As further shown in Figure 4, "[a]pplications exist and run in user mode while the operating system kernel **46** itself runs in kernel mode." '058 Patent at 8:18-20. The '058 Patent explains that "[u]ser mode functionality includes the user applications **42**, the standard application libraries **44**, and one or more critical system elements, **45 & 47**." '058 Patent at 8:18-22. Consistent therewith, the '058 Patent explains that a "user process" includes "the application itself, the regular application library, the extended library, and the critical system element all of which are operating in user mode," allowing a CSE to "execute in the same context as an application." '058 Patent at 1:48, 9:44-47, 9:41-42. Thus, "critical system elements [] exist in the same context as an application." '058 Patent at 9:41-42.

29. Both claim terms addressed herein are found in claim 1 of the '058 Patent, which recites:

1. A computing system for executing a plurality of software applications comprising:

a) a processor;

b) an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor; and,

c) a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode and

i) wherein some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications,

ii) wherein an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function, and

iii) wherein a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.

V. PRIORITY DATES OF THE '058 PATENT

30. I understand that Plaintiff contends that the claims of the '058 Patent are entitled to the priority dates outlined in the table below.

'058 Patent Claims	Priority Date Claimed by Plaintiff
Claims 1-3 and 18	June 18, 2002
Claim 4	July 2, 2002

31. I do not offer an opinion on whether Plaintiff is entitled to these dates, but in forming my opinions, I have applied them in considering the '058 Patent from the perspective of a POSITA, informed by the specification and teachings of the '058 Patent, the intrinsic record, and extrinsic evidence where appropriate, consistent with the legal standards noted above. I further understand that the provisional application for the '058 Patent (No. 60/504,213) was filed on

September 22, 2003. My opinions and analysis of the '058 Patent provided in this Declaration would not change if this later priority date were applied.

VI. CLAIM CONSTRUCTION OPINIONS

A. '058 Patent Terms

1. "critical system elements" / "operating system critical system elements" / "shared library critical system elements"

IBM/HPE: Indefinite	**VirtaMove**: "any service or part of a service, 'normally' supplied by an operating system, that is critical to the operation of software application.

32. Claim 1 of the '058 patent recites "an operating system kernel having OS critical system elements" and "a shared library having shared library critical system elements."

33. In my opinion, the term "critical system elements" had no established meaning in the art at the time of the alleged invention, and a POSITA would not have reasonable certainty about the scope of the term as used in the specification. As such, the term "critical system elements" is indefinite.

34. Although the '058 Patent defines a "critical system element" as "[a]ny service or part of a service, 'normally' supplied by an operating system, that is critical to the operation of a software application," and describes a CSE as "a dynamic object providing some function that is executing instructions used by applications," ('058 Patent at 6:6-10), this explanation does not provide a POSITA with reasonable certainty as to what "system elements" are "critical."

35. First, the specification's explanation that a CSE includes "any service or part of a service, 'normally' supplied by an operating system that is ***critical*** to the operation of a software application," is both subjective and circular, as it essentially defines the term "critical system elements" to be services that are "critical." Whether a service is "critical" to the operation of a software application is subjective, and will depend on the specific context of the software

application, including its specific functionalities. This is further complicated by the fact that a software application could perform functions as if it had a particular specific service through alternative mechanisms. For example, an application can display information using the operating system service, which transfers information to the graphics card connected to a display. The same information can be transmitted over the network using the operating system's networking services and presented to the user through another computing device's display. Given the alternate mechanisms, it is unclear whether either service is "critical" because the software application could perform the same function (i.e., transmitting and displaying information) with either of them. Thus, the universe of considerations for what makes one service "critical" to a software application is boundless.

36. Determining whether a service is "critical" to the operation of a software application also depends on the operating system in question. For example, a service could be "critical" for a software application operating on a Linux-based OS distribution but not critical for a software application operating on a Windows-based OS distribution. When researchers tested the robustness of six Windows variants and the Linux operating system using applications that use individual operating system services one at a time through system calls, they found that what constitutes a critical service depends on the operating system. Functionally similar services offered by the operating systems caused different degrees of failures, from catastrophic (i.e. the operating system seized to function) to non-significant (i.e. the application was able to operate without loss of functionality). *See, e.g.*, Appendix C, Charles P. Shelton, Phillip Koopman & Kobey Devale, *Robustness Testing of the Microsoft Win32 API* (Proceeding International Conference on Dependable Systems and Networks, 2000).

37. To further complicate the issue, an operating system service may be critical to the functioning of the computer as a whole, even though no specific application requires the service to function. For example, the Windows device management services for removable media can cause the operating system to "freeze" when a new device is connected or removed, regardless of whether this new device is being utilized by any application at that time.

38. Thus, the "critical" quality of any given service would not only depend on the software applications a POSITA may consider, but also on the operating systems, the universe of which is not defined by the '058 Patent. Again, the specification and prosecution history of the '058 Patent provide no objective standard to guide a POSITA's determination of what makes a service "critical."

39. Second, a POSITA will understand "normal" to mean "typical" or "expected." But whether a service is "'normally' supplied" by an operating system necessarily differs from one operating system to another. In my experience, different operating systems provide different services for the variety of software applications that run on the operating systems. At the time of the invention, there were many operating systems available with very diverse designs, *e.g.*: MS/PC DOS, Macintosh System 7, Windows 3, Windows 9, Windows 2000, Windows NT, AmigaDOS, hline MTS, VMS, BeOS, and many versions of the Berkeley Unix operating system. In addition, there were many Unix operating systems available from a variety of vendors offering different services and capabilities, including FreeBSD, BSD by Berkeley, SunOS/Solaris by Sun Microsystems, Ultrix and OSF 1/Digital Unix by DEC/Compaq, HPUX by Hewlett-Packard, AIX by IBM, IRIX by Silicon Graphics, GNU/Linux that was open source, and SCO Unix by Novell. These operating systems were further separated between single-user and multi-user. At the same time, some of the operating systems were designed to perform a single task and others were

architected to allow multi-tasking. As such, a service that is said to be "'normally' supplied" by an operating system will necessarily depend on the specifics of the underlying operating systems, the software applications being run, and the services provided therewith.

40. A POSITA would not have been aware of any objective standards to determine whether a "service or parts of a service" was "'normally' supplied" without knowing the particular version of the operating system in question. Though some services are common across operating systems, the '058 Patent provides no objective boundaries with respect to how common a service must be in operating systems before it is considered "'normally' supplied."

41. Further, even for services that were common among some operating systems, the specific architectural implementation and mechanisms through which these services were offered to applications differed in many cases. For example, in the Exokernel Operating System and other libOSes, some of the operating system functionality and services are provided by applications, not the operating system, as depicted in Figure 1 below. Appendix D, Dawson R. Engler, The Exokernel Operating System Architecture, 22 (1998) (Ph.D. dissertation, Massachusetts Institute of Technology).

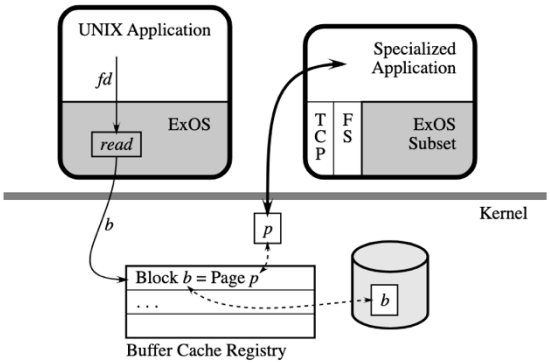

Figure 2-1: A simplified exokernel system with two applications, each linked with its own libOS and sharing pages through a buffer cache registry.

Figure 1: A simplified version of the Exokernel with two applications that use different OS services and functionality to perform the same task (sharing of pages).

42. In addition, some services were more commonly supplied on one operating system than another, given that the universe of services often differed across operating systems and, in some cases, versions of the same operating system. Moreover, distributed and network operating systems may be implemented to serve multiple computing devices and applications, each offering a different set of services.

43. U.S. Patent No. 7,519,814 (the "'814 Patent"), based on provisional application No. 60/512,213, which is incorporated in the '058 Patent, adds further ambiguity to what services can be considered "'normally' supplied by an operating system" insofar as it explains that "containers are created using, for example, ***Linux, Windows, and Unix systems*** . . . however, ***the invention is not limited to these operating systems***." '814 Patent at 7:34-38 (emphasis added). Thus, the '058 Patent acknowledges that it is not limited to any specific set of operating systems. But the specification fails to define the universe of operating systems a POSITA may consider when determining whether a particular service is "'normally' supplied."

44. Moreover, the specification refers to CSEs that are *normally found* or *normally embodied* in an operating system. *See, e.g.*, '058 Patent at 5:21-23 ("[e]mbodiments of the invention enable the replication of critical system elements *normally found* in an operating system kernel") (emphasis added); '058 Patent at 4:55-57 ("In accordance with this invention, critical system elements *normally embodied* in an operating system are exported to software applications through shared libraries") (emphasis added). But there are no objective boundaries to define when services are to be considered "normally" supplied, found, or embodied by any particular operating system.

45. Further, the specification purports to provide examples of "critical system elements," such as:

- TCP/IP, Bluetooth, ATM; or message passing protocols;
- File System services that offer extensions to those supplied by the OS;
- Implementation of file system optimizations for specific application behavior; and
- Implementation of network optimizations for specific application services.

'058 Patent at 6:11-28. However, neither the specification nor the prosecution history explains whether any of the purported examples are "'normally supplied'" by any particular operating system. Additionally, not all software applications require the services outlined in the specification in order to be fully functional. For example, a desktop user using a computer-aided design (CAD) application might not depend on or require Bluetooth and TCP/IP functionality; however, it does depend on the services that offer graphical user interface support. Conversely, a computer running an operating system that runs a web server application heavily depends on TCP/IP and networking services but does not make use of the graphical user interface services provided by the operating system.

46. There exist many other examples depending on the application functionality, the design of the operating system, and the hardware it is executed on. Notwithstanding the examples, a POSITA would not be able to determine if a service is "critical" without objective boundaries as to what specific features or functionalities make the broad categories of listed services "critical." The specification fails to articulate why these examples are purportedly "critical" to any particular software application. As such, even if relying on the broad examples in the specification, a POSITA would still be unable to determine what constitutes a "critical" service."

47. Extrinsic evidence further confirms my opinion that the term "critical" is subjective, having no established meaning in the context of software and computer engineering. I consulted the IEEE Dictionary, published in 2000, which provides technical definitions for terms used in the context of computer engineering. Although the IEEE Dictionary does not provide a

standalone definition for "critical" (in the context of application software or services), it defines several terms which incorporate the word. For example, the IEEE Dictionary defines "criticality" as: "A ***subjective*** description of the intended use and application of the system. Software criticality properties ***may include safety, security, complexity, reliability, performance, or other characteristics.***" *Criticality*, IEEE 100, THE AUTHORITATIVE DICTIONARY OF IEEE STANDARDS TERMS (7th ed. 2000) (emphasis added) (IBMHPE_VM_000000018). This definition confirms that "criticality" in the software context is a subjective determination. For a service to be deemed "critical" depends on several unenumerated properties, including those not explicitly listed in the definition of the term.

48. As another example, the IEEE Dictionary defines "criticality" as "[t]he ***degree*** of impact that a requirement, module, error, fault, failure, or other item has on the development or operation of a system. *Synonym:* severity." *Criticality*, IEEE 100, THE AUTHORITATIVE DICTIONARY OF IEEE STANDARDS TERMS (7th ed. 2000) (emphasis added) (IBMHPE_VM_000000018). And as another example, the IEEE Dictionary defines "critical piece first (software)" as: "A system development approach in which the most critical aspects of a system are implemented first. The critical piece ***may be defined*** in terms of ***services provided, degree of risk, difficulty, or other criteria***." *Critical piece first*, IEEE 100, THE AUTHORITATIVE DICTIONARY OF IEEE STANDARDS TERMS (7th ed. 2000) (emphasis added) (IBMHPE_VM_000000019). These definitions further reinforce the subjectivity of the term "critical," indicating that it is a term of degree. Although different services may have varying degrees of criticality to a particular software application, the '058 Patent fails to provide objective boundaries for a POSITA to discern the degree of criticality required for a "system element" to be "critical" to any particular software application.

49. Because the '058 Patent fails to provide any objective standard for determining what is "'normally' supplied" and "critical" as used in the definition of "critical system elements," a POSITA would not have reasonable certainty about the scope of term "critical system elements." Therefore, in my opinion, asserted claims 1, 2, 3, 4, and 18 of the '058 Patent reciting "critical system elements" are indefinite.

2. "functional replicas"

IBM/HPE: Indefinite	**VirtaMove**: "substantial functional equivalents or replacements of kernel functions"

50. Claim 1 also requires that the recited critical system elements (CSE) in the shared library be "functional replicas" of the CSEs in the operating system ("OSCSE").

51. Based on my experience, the term "functional replicas" had no established meaning in the art at the time of the alleged invention and in my opinion, this term is also indefinite.

52. A POSITA would not have reasonable certainty about the scope of "functional replicas." The '058 Patent provides the following definition for "replica": "The term replica used herein is meant to denote a CSE having similar attributes to, but not necessarily and preferably not an exact copy of a CSE in the operating system (OS); notwithstanding, a CSE for use in user mode, may in a less preferred embodiment be a copy of a CSE in the OS." '058 Patent at 1:66-2:3.

53. The definition uses subjective language by requiring that a "replica" have "*similar attributes* to" a CSE in the operating system. But, whether a CSE has "similar attributes" is a subjective determination. For example, CSEs may be said to be "similar" by virtue of their underlying application codes; by having the same application programming interfaces, or by accomplishing or because they accomplish the same result, but with different backend operations or dependencies to other existing services. There are simply no objective bounds provided by the specification regarding what constitutes "similar attributes." Instead, the '058 Patent adds further

vagueness by noting that a replica is "not *necessarily* and preferably not an exact copy." '058 Patent at 1:67-2:1 (emphasis added). Neither the specification nor the prosecution history of the '058 Patent provides any guidance as to how "similar" a CSE must be to constitute a replica of a CSE in the operating system, much less what properties of a CSE a POSITA should consider in determining whether a CSE is sufficiently similar to another CSE in the operating system, without being an exact copy.

54. Indeed, the patent's description of the preferred embodiment of the invention invites additional confusion. For example, the specification explains that "[a] CSE is replicated by way of a kernel service being repeated in user space. ***This replicated CSE may differ slightly from its counterpart in the OS***. Replication is achieved placing ***CSEs similar to those in the OS in shared libraries*** which provides a means of attaching or linking a CSE service to an application having access to the shared library. ***Therefore, a service in the kernel is substantially replicated in user mode through the use of shared libraries***." '058 Patent at 5:20-34 (emphasis added).

55. Here again, the '058 Patent explains that a CSE is "replicated by way of a kernel service being repeated in user space." That a CSE is "replicated" by being "***repeated***" would suggest to a POSITA that the CSE is an "exact copy of a CSE in the operating system." But this cannot be the case because the definition for "replica" discourages having "an exact copy." 058 Patent at 1:66-2:3. The specification further explains that "[t]his replicated CSE *may differ slightly* from its counterpart in the OS." '058 Patent at 5:28-30 (emphasis added). But this purported explanation introduces yet another term of degree in the form of "slightly." Nothing in the specification or the prosecution history provides objective boundaries as to what would constitute a "difference," or the degree to which the differences between a CSE and a CSE in the operating system are sufficiently "slight" to be deemed "similar."

56. Further, moving code from the kernel to the user space of an application would inherently require a substantially different design architecture and functional implementation, as the code would no longer be part of the operating system and have direct access to kernel interfaces, modules, and functionality. The code design for kernel functions substantially differs from that of user space services, even when offering the same functionality, and has a different set of functional and code dependencies.

57. Further, the specification explains that "[r]eplication is achieved by placing CSEs similar to those in the OS in shared libraries." '058 Patent at 5:28-30. This also leaves open the question of what makes CSEs in shared libraries "similar" to those in the OS. Also vague is the specification's explanation that "a service in the kernel is substantially replicated in user mode through the use of shared libraries." *Id.* Here, the specification introduces additional terms of degree and uncertainty, like "substantially replicated."

58. Further, the specification explains that:

> The CSE library includes ***replicas*** or ***substantial functional equivalents*** or ***replacements of kernel functions***. ***The term replica, shall encompass any of these meanings, and although not a preferred embodiment, may even be a copy of a CSE that is part of the OS.*** These functions can be directly called by the applications **42** and as such can be run in the same context as the applications **42**. In preferred embodiments, the kernel functions which are included in the extended shared library **45,47** and critical system element library are also included in the operating system kernel **46**.

'058 Patent at 8:27-44 (emphasis added). First, the statement defines "replica" as including any of these following meanings: "replicas," "substantial functional equivalents," and "replacements of kernel functions." Defining the term "replica" with itself does not provide any useful guidance to a POSITA. Moreover, in my opinion, the other supposedly definitional terms, including "substantial functional equivalent," are similarly imprecise and ambiguous.

59. With respect to "substantial functional equivalent," the term "substantial" is a subjective term of degree. Neither the specification nor the file history provide any guidance as to what "substantial functional" entails and how a POSITA would measure how "substantial" the "functional equivalence" must be for a CSE to be considered a "replica" of a CSE in the operating system.

60. Further, in describing preferred embodiments of the invention, the specification explains that:

> In preferred embodiments, the critical system elements which are included in user mode are replicas of elements which are still included in the operating system kernel. The term replication means that ***like services*** are supplied. As was described heretofore, it is not necessarily the case that duplicates of the same implementation found in the kernel are provided by a CSE; but ***essentially the same functionality*** is provided.

'058 Patent at 9:41-58 (emphasis added). This statement repeats the same ambiguity concerning the meaning of "replica." Specifically, the statement replaces the term "similar to" with synonyms in the form of "like services" and "essentially the same functionality." *See, e.g.*, *id.*; '058 Patent at 5:28-30. A POSITA would not have any objective boundaries to determine what constitutes a replica supplying a "like service" and/or "essentially the same functionality," much like there are no objective boundaries in the specification to decipher what constitutes a CSE "similar to" a CSE in an operating system.

61. Further, the term "***functional*** replicas" is only used in the claim (and nowhere else in the specification). As discussed above, the specification offers definitions for "replica" and "replication." Adding the word "functional" before "replica" does not address the ambiguity created by the specification's subjective definition for "replica." The only other use of "functional replica" occurred when the patentee sought to distinguish a prior art reference, explaining that "[n]owhere does Cabrero et al. disclose ***the SLCSEs stored in the shared library being functional***

replicas of OSCSEs, or in other words, replacements." '058 Patent File History, July 1, 2009 Response at 8 (emphasis added). However, even if a replica is a "replacement" of a CSE in the operating system, a POSITA would still understand that this "replacement" must be "similar" to the CSE in the operating system, given that it is "preferably not" an "exact copy of a CSE in the operating system." '058 Patent at 1:66-2:3, 5:28-30.

62. Extrinsic evidence further confirms my opinion that the term "functional replica" is unclear and ambiguous when read in light of the specification. In the context of software, the IEEE Dictionary defines "function" as: "(A) A ***defined objective*** or ***characteristic action*** of a system or component. For example, a system may have inventory control as its primary function. (B) A software module that performs a ***specific action***, is invoked by the appearance of its name in an expression, may receive input values, and returns a single value." *Function*, IEEE 100, THE AUTHORITATIVE DICTIONARY OF IEEE STANDARDS TERMS (7th ed. 2000) (emphasis added) (IBMHPE_VM_000000010). The IEEE Dictionary also defines "function" as: "(8) (software user documentation) A ***specific*** purpose of an entity or it characteristic action." *Id.* (emphasis added). And Webster's II New College Dictionary (2001) provides an ordinary definition for "replica" as: "A ***copy*** or ***reproduction***." *Replica*, WEBSTER'S II NEW COLLEGE DICTIONARY (2001) (emphasis added) (IBMHPE_VM_000000025).

63. Based on these extrinsic evidence, a POSITA may understand "functional replica" of a CSE in an operating system to be, "a copy of a CSE performing the same specific action (or defined objective) as the CSE in the operating system." Yet, the patent specification obviates any objective boundaries for the term "replica" by using terms such as "similar," "not an exact copy," "may even be a copy," "substantial functional equivalent," "like," and "essentially the same functionality."

64. Because the ’058 Patent fails to provide any objective standard for determining the meaning of the term “functional replicas,” a POSITA would not have reasonable certainty about the scope of this term. Therefore, in my opinion, asserted claim 1 of the ’058 Patent reciting “functional replicas” is indefinite.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct. Executed this 10th day of January, 2025.

/s/ Angelos Stavrou

Angelos Stavrou

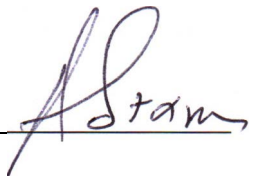

APPENDIX A

Materials Considered

1. VirtaMove's Asserted Patents

   U.S. Patent No. 7,784,058

   U.S. Patent No. 7,519,814

2. File Histories of the Asserted Patents

3. Extrinsic Evidence

   Criticality, IEEE 100, The Authoritative Dictionary of IEEE Standards Terms (7th ed. 2000)

   Critical piece first, IEEE 100, The Authoritative Dictionary of IEEE Standards Terms (7th ed. 2000)

   Function, IEEE 100, The Authoritative Dictionary of IEEE Standards Terms (7th ed. 2000)

   Replica, Webster's II New College Dictionary (2001)

   Charles P. Shelton, Phillip Koopman & Kobey Devale, *Robustness Testing of the Microsoft Win32 API* (Proceeding International Conference on Dependable Systems and Networks, 2000).

   Dawson R. Engler, The Exokernel Operating System Architecture, 22 (1998) (Ph.D. dissertation, Massachusetts Institute of Technology).

4. Correspondences Between Counsel for the Parties Regarding Proposed Constructions

5. The Parties' EDTX Local Rule 4 Disclosures

APPENDIX B

ANGELOS STAVROU - *Curriculum Vitae*

POSITIONS HELD

- **November 2022 – Present**
  Founder & CEO, Impedyme Inc., Arlington, VA
- **March 2022 – Present**
  Founder & Chairman of the Board, Quokka Inc, Arlington, VA
- **March 2011 – Present**
  Founder & CEO, Kryptowire LLC, Arlington, VA
- **August 2020 – Present**
  Entrepreneurship Lead, Innovation Campus, Virginia Tech, Arlington, VA
- **August 2020 – Present**
  Professor, Innovation Campus & Bradley Department of Electrical and Computer Engineering, Virginia Tech, Arlington, VA
- **August 2017 – August 2020**
  Professor, Computer Science Department, George Mason University, Fairfax, VA
- **April 2014 – August 2020**
  Director, Center for Assurance Research and Engineering (CARE), George Mason University, Fairfax, VA
- **May 2012 – August 2017**
  Associate Professor, Computer Science Department, George Mason University, Fairfax, VA
- **March 2014 – May 2017**
  Academic Director, M.S. in Management of Secure Information Systems Program, School of Management, George Mason University, Fairfax, VA
- **May 2013 – August 2015**
  Academic Director, M.S. in Information Security and Assurance, Computer Science Department, George Mason University, Fairfax, VA
- **September 2011 – August 2020**
  Associate Researcher, National Institute of Standards and Technology (NIST), Computer Security Division, Gaithersburg, MD
- **January 2012 – August 2013**
  Associate Director, Center for Secure Information Systems, George Mason University, Fairfax, VA
- **August 2007 – May 2012**
  Assistant Professor, Computer Science Department, George Mason University, Fairfax, VA
- **May 2006 – August 2006**
  Research Intern, Microsoft Research, Cambridge, UK
- **August 2004 – December 2004**
  Software Engineer Intern, Google Inc., Mountain view, CA
- **September 2001 – August 2007**
  Research Assistant, Computer Science Department, Columbia University, New York, NY

EDUCATION

- **Columbia University, Fu Foundation School of Engineering & Applied Science, New York, NY.**
  Ph.D. in Computer Science **(with Distinction)** (**August 2007**)
  Thesis: "An Overlay Architecture for End-to-End Service Availability".
  Advisor: Angelos D. Keromytis.

- **Columbia University, Fu Foundation School of Engineering & Applied Science, New York, NY.**
  M.Phil. Degree in Computer Science (**January 2007**)
  M.Sc. Degree in Electrical Engineering with concentration in Multimedia Networking (Peer to Peer Networks). (December 2002).

- **National University of Athens, Athens Greece /Carleton University ON, Canada.**
  M.Sc. Degree in Algorithms, Computability and Logic. (**June 2001**)
  Master's Thesis: "A new distributed algorithm for routing in satellite constellation networks" Advisor: Prof. E. Kranakis.

- **University of Patras, Electrical Engineering Department, Patras Greece.**
  Certificate of Engineering for the completion of the last two years of coursework in Electrical Engineering (**February 1999**).

- **University of Patras, Physics Department, Patras, Greece.**
  B.Sc (Honors) in Physics, (**July 1997**) Thesis: "Stream Ciphers theory and practical application".

Publications

Issued Patents

1. **Systems and methods for analyzing software**
   Ryan Johnson, Nikolaos Kiourtis, Angelos Stavrou
   U.S. Patent Number 10,387,627. Issued on August 20th, 2019.

2. **Active Authentication of Users**
   Angelos Stavrou, Rahul Murmuria, Ryan Johnson, Daniel Barbara
   U.S. Patent Number 10,289,819. Issued on May 14th, 2019.

3. **Methods and systems for increased debugging transparency**
   Fengwei Zhang, Kevin Leach, Angelos Stavrou, Haining Wang
   U.S. Patent Number 10,127,137. Issued on November 13th, 2018.

4. **Methods and Apparatus for Application Isolation**

Anup Ghosh, Yih Huang, Jiang Wang, Angelos Stavrou
U.S. Patent Number 9,602,524. Issued on March 21st, 2017.

5. **Malware Detector**
Angelos Stavrou, Sushil Jajodia, Anup Ghosh, Rhandi Martin, Charalampos Andrianakis
U.S. Patent Number 9,531,747. Issued on December 27th, 2016.

6. **Hardware-assisted Integrity Monitor**
Jiang Wang, Anup Ghosh, Kun Sun, Angelos Stavrou
U.S. Patent Number 9,270,697. Issued on February 23rd, 2016.

7. **Systems and Methods for Inhibiting Attacks with a Network**
Angelos Stavrou, Angelos D. Keromytis
U.S. Patent Number 9,344,418. Issued on May 17th, 2016.

8. **Methods and Apparatus for Application Isolation**
Anup Ghosh, Yih Huang, Jiang Wang, Angelos Stavrou
U.S. Patent Number 9,098,698. Issued on August 4th, 2015.

9. **Adaptive feedback loop based on a sensor for streaming static and interactive media content to animals**
Angelos Stavrou, Margaret Lee Perry-Flippin
U.S. Patent Number 9,043,818. Issued on May 26th, 2015.

10. **Malware Detector**
Angelos Stavrou, Sushil Jajodia, Anup Ghosh, Rhandi Martin, Charalampos Andrianakis
U.S. Patent Number 8,935,773. Issued on January 13th, 2015.

11. **Systems, Methods, and Media for Recovering an Application from a Fault or Attack**
Michael E. Locasto, Angelos D. Keromytis, Angelos Stavrou, Gabriela F. Ciocarlie
U.S. Patent Number 8,924,782. Issued on December 30th, 2014.

12. **Hardware-assisted Integrity Monitor**
Jiang Wang, Angelos Stavrou, Anup Ghosh, Kun Sun
U.S. Patent Number 8,819,225. Issued on August 26th, 2014.

13. **Website Matching based on Network Traffic**
Angelos Stavrou, Mohammed A. Alhussein, Brian Sanders
U.S. Patent Number 8,726,005. Issued on May 13th, 2014.

14. **Systems and Methods for Inhibiting Attacks with a Network**
Angelos Stavrou, Angelos D. Keromytis.
U.S. Patent Number 8,631,484. Issued on January 14th, 2014.

15. **Methods, Media and Systems for Responding to a Denial of Service Attack**
Angelos Stavrou, Angelos D. Keromytis, Jason Nieh, Vishal Misra, and Daniel Rubenstein.
U.S. Patent Number 8,549,646. Issued on October 1st, 2013.

16. **Systems, Methods, and Media for Generating Sanitized Data, Sanitizing Anomaly Detection Models, and/or Generating Sanitized Anomaly Detection Models**
Gabriela Cretu, Angelos Stavrou, Salvatore J. Stolfo, Angelos D. Keromytis, Michael E. Locasto.
U.S. Patent Number 8,407,160. Issued on March 26th, 2013.

17. **Systems and Methods for Computing Data Transmission Characteristics of a Network Path Based on Single-ended Measurements**
Angelos D. Keromytis, Sambuddho Chakravarty, and Angelos Stavrou.
U.S. Patent Number 8,228,815. Issued on July 24th, 2012.

18. **Methods, Systems and Media for Software Self-Healing**
Michael E. Locasto, Angelos D. Keromytis, Salvatore J. Stolfo, Angelos Stavrou, Gabriela Cretu,
Stylianos Sidiroglou, Jason Nieh, and Oren Laadan.
U.S. Patent Number 7,962,798. Issued on June 14th, 2011.

19. **Systems and Methods for Computing Data Transmission Characteristics of a Network Path Based on Single-ended Measurements**
Angelos D. Keromytis, Sambuddho Chakravarty, and Angelos Stavrou. U.S. Patent Number 7,660,261. Issued on February 9th, 2010.

Journal Publications

1. **Revealing Protocol Architecture's Design Patterns in the Volumetric DDoS Defense Design Space**
Zhang, Zhiyi; Xiao, Guorui; Song, Sichen; Aygun, R. Can; Stavrou, Angelos; Zhang, Lixia; Osterweil, Eric. To appear in the proceedings of IEEE Communications Surveys and Tutorials (2024).

2. **A multiview clustering framework for detecting deceptive reviews**
Yubao Zhang, Haining Wang, Angelos Stavrou. In the proceedings of Journal of Computer Security, vol. 32, no. 1, pp. 31-52, 2024.

3. **OFDRA: Optimal Femtocell Deployment for Accurate Indoor Positioning of RIS-Mounted AVs**
A. Famili, T. O. Atalay, A. Stavrou, H. Wang and J. -M. Park. In the proceedings of IEEE Journal on Selected Areas in Communications, vol. 41, no. 12, pp. 3783-3798, Dec. 2023, doi:10.1109/JSAC.2023.3322821.

4. **iDROP: Robust Localization for Indoor Navigation of Drones With Optimized Beacon Placement**
A. Famili, A. Stavrou, H. Wang and J. -M. Park. In the proceedings of IEEE Internet of Things Journal, vol. 10, no. 16, pp. 14226-14238, 15 Aug.15, 2023, doi: 10.1109/JIOT.2023.3280084.

5. **Revisiting the Spaceborne Illuminators of Opportunity for Airborne Object Tracking. [PDF]** John Robie, Alireza Famili, Angelos Stavrou. In the proceedings of IEEE Computer 56(1): 82-92 (2023)

6. **PILOT: High-Precision Indoor Localization for Autonomous Drones**
A. Famili, A. Stavrou, H. Wang and J. -M. Park. In the proceedings of IEEE Transactions on Vehicular Technology, vol. 72, no. 5, pp. 6445-6459, May 2023, doi: 10.1109/TVT.2022.3229628

7. **Understanding the Security Implication of Aborting Live Migration**
X. Gao, J. Xiao, H. Wang and Angelos Stavrou. IEEE Trans. Cloud Comput. 10(2)**:** *1275-1286 (Online: 2020, Published: 2022)*

8. **21 Years of Distributed Denial-of-Service: A Call to Action – Part 2 [PDF]**
Eric Osterweil, Angelos Stavrou, and Lixia Zhang. In *Computer*, vol. 53, no. 8, pp. 94-99, Aug. 2020, doi: 10.1109/MC.2020.2993330.

9. **21 Years of Distributed Denial-of Service: Current State of Affairs - Part 1 [PDF]**
Eric Osterweil, Angelos Stavrou, and Lixia Zhang. In Computer, vol. 53, no. 7, pp. 88-92, July 2020, doi: 10.1109/MC.2020.2983711.

10. **Towards Transparent Debugging**
Fengwei Zhang, Kevin Leach, Angelos Stavrou, Haining Wang
In the proeedings of IEEE Transactions of Dependable Secure Computing 15(2): 321-335 (IEEE TDSC 2018)

11. **On Early Detection of Application-level Resource Exhaustion and Starvation**
Mohamed Elsabagh, Daniel Barbará, Dan Fleck, Angelos Stavrou
In proceedings of the Elsevier Journal of Systems and Software 137: 430-447 (2018)

12. **An Empirical Investigation of Ecommerce-Reputation-Escalation-as-a-Service**
Haitao Xu, Daiping Liu, Haining Wang, Angelos Stavrou.
In the proceedings of ACM Transactions on the Web, Volume 11, Number 2, May 2017 (pages 13:1-13:35)

13. **DDoS in the IoT: Mirai and Other Botnets**
Constantinos Kolias, Georgios Kambourakis, Angelos Stavrou, Jeffrey M. Voas.

In the proceedings of IEEE Computer 50(7): 80-84 (2017)

14. **Cybersecurity Leadership: Competencies, Governance, and Technologies for Industrial Control**
Jean-Pierre Auffret, Jane L. Snowdon, Angelos Stavrou, Jeffrey S. Katz, Diana Kelley, Rasheq S. Rahman, Frank Stein, Lisa Sokol, Peter Allor, Peng Warweg. Systems. Journal of Interconnection Networks 17(1): 1-20 (2017)

15. **Verified Time**
Angelos Stavrou, Jeff Voas.
In the proceedings of IEEE Computer, Volume: 50, Issue: 3, March 2017

16. **On the Move: Evading Distributed Denial-of-Service Attacks**
Angelos Stavrou, Daniel Fleck, Constantinos Kolias.
In the proceedings of IEEE Computer 49(3): 104-107 (2016)

17. **Intrusion Detection in 802.11 Networks: Empirical Evaluation of Threats and a Public Dataset**
Constantinos Kolias, Georgios Kambourakis, Angelos Stavrou, Stefanos Gritzalis.
In the proceedings of IEEE Communications Surveys and Tutorials 18(1): 184-208 (2016)

18. **Learning Internet-of-Things Security "Hands-On"**
Constantinos Kolias, Angelos Stavrou, Jeffrey M. Voas, Irena Bojanova, D. Richard Kuhn.
In the proceedings of IEEE Security & Privacy 14(1): 37-46 (2016)

19. **Securely Making "Things" Right**
Constantinos Kolias, Angelos Stavrou, Jeffrey M. Voas.
In the proceedings of IEEE Computer 48(9): 84-88 IEEE Magazine (2015)

20. **A Moving Target DDoS Defense Mechanism**
Huangxin Wang, Quan Jia, Dan Fleck, Walter Powell, Fei Li, Angelos Stavrou.
In the proceedings of Elsevier Journal of Computer Communications, 46: 10-21 (2014)

21. **HyperCheck: A Hardware-Assisted Integrity Monitor**
Fengwei Zhang, Jiang Wang, Kun Sun, and Angelos Stavrou.
In the proceedings of IEEE Transactions on Dependable and Secure Computing (TDSC), 11(4): 332-344 (2014)

22. **Improving network response times using social information**
Sharath Hiremagalore, Chen Liang, Angelos Stavrou and Huzefa Rangwala.
In the proceedings of Social Network Analysis and Mining, Springer
Social Network Analysis and Mining, Volume 3, pages 209-220 (2013)

23. **Providing Users' Anonymity in Mobile Hybrid Networks**
Claudio Agostino Ardagna, Sushil Jajodia, Pierangela Samarati, Angelos Stavrou.
In the proceedings of ACM Transactions on Internet Technology, Volume 12, 3, Article 7, pages 1 - 33 (May 2013)

24. **Building Security into Off-the-Shelf Smartphones**
Angelos Stavrou, Jeffrey Voas, Tom Karygiannis, Steve Quirolgico.
In the proceedings of IEEE Computer, vol. 45, no. 2, pp. 82-84, Feb. 2012, doi:10.1109/MC.2012.44

25. **DoubleGuard: Detecting Intrusions In Multi-tier Web Applications**
Meixing Le, Angelos Stavrou, Brent ByungHoon Kang.
In the proceedings of IEEE Journal on Transactions on Dependable and Secure Computing (TDSC) 2011, ISSN: 1545-5971 10 Nov. 2011. IEEE computer Society Digital Library. IEEE Computer Society. Acceptance Rate: 10-12% as reported by 2009 TDSC editorial, ISI Impact Factor: 2.093 (2010).

26. **The Ephemeral Legion: Producing an Expert Cyber-security Workforce from Thin Air**
Michael E. Locasto, Anup Ghosh, Sushil Jajodia, and Angelos Stavrou.
In the proceedings of Communications of the ACM, Vol. 54, Issue 1, pp 129 - 131. Impact Factor: 2.362 (2010). [bib]

27. **The Dynamic Community of Interest and its Realization in ZODIAC**
Scott Alexander, Steve Bellovin, Yuu-Heng Cheng, Brian Coan, Andrei Ghetie, Vikram Kaul, Nicholas F. Maxemchuk, Henning Schulzrinne, Stephen Schwab, Bruce Siegell, Angelos Stavrou, and Jonathan M. Smith.
In the proceedings of IEEE Communications Magazine, October 2009, pp. 40-47. Impact Factor: 2.837

28. **On the Infeasibility of Modeling Polymorphic Shellcode: Re-thinking the Role of Learning in Intrusion Detection Systems**
Yingbo Song, Michael E. Locasto, Angelos Stavrou, Angelos D. Keromytis, and Salvatore J. Stolfo.
In the Proceedings of Machine Learning Journal (MLJ) p. 179-205. Accepted: 7 August 2009.
Published online: 29 October 2009. Editors: Pavel Laskov and Richard Lippmann. ISI Impact Factor: 1.956 (2010). [bib]

29. **WebSOS: An Overlay-based System For Protecting Web Servers From Denial of Service Attacks**
Angelos Stavrou, Debra L. Cook, William G. Morein, Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein.
In the proceedings of Elsevier Journal of Computer Networks, special issue on Web and Network Security, vol. 48, no.5, p. 781 - 807. August 2005 5-Year Impact Factor: 1.690. [bib]

30. **A Lightweight, Robust, P2P System to Handle Flash Crowds**
Angelos Stavrou, Dan Rubenstein, Sambit Sahu.
In the Proceedings of IEEE Journal on Selected Areas in Communications (JSAC), special issue on Service Overlay Networks, Volume 22, Number 1, p. 6-17, January 2004. Impact Factor: 4.232 (2010). [bib]

Conference Publications

1. **2023. Dial "N" for NXDomain: The Scale, Origin, and Security Implications of DNS Queries to Non-Existent Domains.**
Guannan Liu, Lin Jin, Shuai Hao, Yubao Zhang, Daiping Liu, Angelos Stavrou, and Haining Wang.
In the proceedings of the 2023 ACM on Internet Measurement Conference (IMC '23). Association for Computing Machinery, New York, NY, USA, 198–212.
https://doi.org/10.1145/3618257.3624805

2. **Securing 5G OpenRAN with a Scalable Authorization Framework for xApps**
Tolga O Atalay, Sudip Maitra, Dragoslav Stojadinovic, Angelos Stavrou, Haining Wang. In the proceedings of IEEE INFOCOM 2023 - IEEE Conference on Computer Communications, New York City, NY, USA, 2023, pp. 1-10, doi: 10.1109/INFOCOM53939.2023.10228961.

3. **Vehicular Teamwork for Better Positioning**
A. Famili, V. Slyusar, Y. H. Lee and A. Stavrou. In the proceedings of 2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Honolulu, Oahu, HI, USA, 2023, pp. 3507-3513, doi: 10.1109/SMC53992.2023.10393920.

4. **EGO-6: Enhancing Geofencing Security Systems with Optimal Deployment of 6G TRPs**
Alireza Famili, Angelos Stavrou, Haining Wang, Jung-Min Jerry Park. In the proceedings of Silicon Valley Cybersecurity Conference, SVCC 2023, San Jose, CA, USA, May 17-19, 2023. *IEEE 2023, ISBN 979-8-3503-2157-9*

5. **Wi-Five: Optimal Placement of Wi-Fi Routers in 5G Networks for Indoor Drone Navigation**
A. Famili, T. Atalay, A. Stavrou and H. Wang. In the proceedings of the 2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring), Florence, Italy, 2023, pp. 1-7, doi: 10.1109/VTC2023-Spring57618.2023.10201144.

6. **Isochrons in tunable photonic oscillators and applications in precise positioning**
G Himona, A Famili, A Stavrou, V Kovanis, Y Kominis
Physics and Simulation of Optoelectronic Devices XXXI 12415, 82-86

7. **Detecting and Measuring Misconfigured Manifests in Android Apps**
Allen Yuqing Yang, Mohamed Elsabagh, Chaoshun Zuo, Ryan Johnson, Angelos Stavrou, Zhiqiang Lin ACM CCS '22: Proceedings of the 2022 ACM SIGSAC

Conference on Computer and Communications SecurityNovember 2022 Pages 3063–3077, https://doi.org/10.1145/3548606.3560607

8. **Network-Slice-as-a-Service Deployment Cost Assessment in an End-to-End 5G Testbed**
   Tolga O. Atalay, Dragoslav Stojadinovic, Alireza Famili, Angelos Stavrou, Haining Wang.
   In the Proceedings of IEEE Global Communications Conference, Rio de Janeiro, Brazil, 2022, pp. 2056-2061, doi: 10.1109/GLOBECOM48099.2022.10001579.

9. **GPS Spoofing Detection by Leveraging 5G Positioning Capabilities**
   Alireza Famili, Mahsa Foruhandeh, Tolga Atalay, Angelos Stavrou, Haining Wang
   IEEE Latin-American Conference on Communications (LATINCOM), Rio de Janeiro, Brazil, 2022, pp. 1-6, doi: 10.1109/LATINCOM56090.2022.10000569.

10. **SPIN: Sensor Placement for Indoor Navigation of Drones**
    Alireza Famili, Angelos Stavrou, Haining Wang, Jung-Min Jerry Park
    In the Proceedings of IEEE Latin-American Conference on Communications (LATINCOM), Rio de Janeiro, Brazil, 2022, pp. 1-6, doi: 10.1109/LATINCOM56090.2022.10000583.

11. **Streaming and Unbalanced PSI from Function Secret Sharing**
    Samuel Dittmer, Yuval Ishai, Steve Lu, Rafail Ostrovsky, Mohamed Elsabagh, Nikolaos Kiourtis, Brian Schulte, Angelos Stavrou. Security and Cryptography for Networks. SCN 2022. Lecture Notes in Computer Science, vol 13409. Springer, Cham. https://doi.org/10.1007/978-3-031-14791-3_25

12. **Eternal flying: Optimal placement of wireless chargers for nonstop drone flights**
    Alireza Famili, Angelos Stavrou
    International Conference on Electrical, Computer and Energy Technologies (ICECET), Prague, Czech Republic, 2022, doi: 10.1109/ICECET55527.2022.9873507.

13. **Receiver Density Analysis for High Probability Detection of Forward Scattered Airborne Signals**
    John Robie, Alireza Famili, Angelos Stavrou
    International Conference on Electrical, Computer and Energy Technologies (ICECET), Prague, Czech Republic, 2022, doi: 10.1109/ICECET55527.2022.9872553

14. **RAIL: Robust Acoustic Indoor Localization for Drones**
    A. Famili, A. Stavrou, H. Wang and J. -M. J. Park
    In the Proceedings of IEEE 95th Vehicular Technology Conference: (VTC2022-Spring), Helsinki, Finland, 2022, pp. 1-6, doi: 10.1109/VTC2022-Spring54318.2022.9860933.

15. **Characterization of AES Implementations on Microprocessor-based IoT Devices**
S. Roy, A. Stavrou, B. L. Mark, K. Zeng, S. M. P D and K. N. Khasawneh.
In the proceedings of IEEE Wireless Communications and Networking Conference (WCNC), Austin, TX, USA, 2022, pp. 55-60, doi: 10.1109/WCNC51071.2022.9771975.

16. **Scaling Network Slices with a 5G Testbed: A Resource Consumption Study**
Tolga O. Atalay, Dragoslav Stojadinovic, Angelos Stavrou, Haining Wang
In the proceedings of IEEE Wireless Communications and Networking Conference (WCNC), Austin, TX, USA, 2022, pp. 2649-2654, doi: 10.1109/WCNC51071.2022.9771860.

17. **DEFInit: An Analysis of Exposed Android Init Routines**
Yuede Ji, Mohamed Elsabagh, Ryan Johnson, and Angelos Stavrou
In the proceedings of the 30th USENIX Security Symposium (USENIX Security 2021)

18. **(Un)protected Broadcasts in Android 9 and 10**
Ryan Johnson, Mohamed Elsabagh, and Angelos Stavrou
BlackHat Asia 2021

19. **Black-Box IoT: Authentication and Distributed Storage of IoT Data from Constrained Sensors**
Panagiotis Chatzigiannis, Foteini Baldimtsi, Constantinos Kolias, and Angelos Stavrou. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation* (*IoTDI 2021*), *2021*

20. **CloudSkulk: A Nested Virtual Machine Based Rootkit and Its Detection.**
J. Connelly, T. Roberts, X. Gao, J. Xiao, H. Wang, and A. Stavrou. *In 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2021), 2021*

21. **FIRMSCOPE: Automatic Uncovering of Privilege-Escalation Vulnerabilities in Pre-Installed Apps in Android Firmware.** **[PDF]**
Mohamed Elsabagh, Ryan Johnson, and Angelos Stavrou, Kryptowire; Chaoshun Zuo, Qingchuan Zhao, and Zhiqiang Lin, The Ohio State University. In the 29th USENIX Security Symposium (USENIX Security 20).

22. **Resilient and Scalable Cloned App Detection Using Forced Execution and Compression Trees**
Mohamed Elsabagh, Ryan Johnson, Angelos Stavrou
In proceedings of the IEEE Conference on Dependable and Secure Computing (DSC 2018)

23. **An adversarial coupon-collector model of asynchronous moving-target defense against botnet reconnaissance**
G Kesidis, Y Shan, D Fleck, A Stavrou, T Konstantopoulos
In proceedings of the 2018 13th IEEE International Conference on Malicious and Unwanted Software (IEEE MALCON)

24. **Moving-target Defense against Botnet Reconnaissance and an Adversarial Coupon-Collection Model**
Dan Fleck, Angelos Stavrou, George Kesidis, N Nasiriani, Y Shan, T Konstantopoulos
In proceedings of the IEEE Conference on Dependable and Secure Computing (DSC 2018)

25. **End Users Get Maneuvered: Empirical Analysis of Redirection Hijacking in Content Delivery Networks**
Shuai Hao, Yubao Zhang and Haining Wang, Angelos Stavrou
In the proceeding of the 27th Usenix Security Symposium, (Usenix Security 2018) August 15-17, 2018, Baltimore, MD, USA

26. **Dazed Droids: A Longitudinal Study of Android Inter-App Vulnerabilities**
Ryan Johnson, Mohamed Elsabagh, Angelos Stavrou, and Jeff Offutt
In the Proceedings of ACM ASIA Conference on Computer & Communications Security 2018,
(ASIACCS 2018), 777-791, June 4 - 8, 2018, Sogdo, Incheon, Korea

27. **Detecting and Characterizing Web Bot Traffic in a Large E-commerce Marketplace**
Haitao Xu, Zhao Li, Chen Chu, Yuanmi Chen, Yifan Yang, Haifeng Lu, Haining Wang, and Angelos Stavrou, In the Proceedings of the 23rd European Symposium on Research in Computer Security (*ESORICS'18*), Barcelona, Spain, Sep. 2018. (Acceptance Rate: 19.8%, 56/283).

28. **The Mirai Botnet and the IoT Zombie Armies**
Georgios Kambourakis, Constantinos Kolias, and Angelos Stavrou
In the Proceedings of the IEEE Military Communications Conference (MILCOM 2017)
October 23 -25, 2017, Baltimore, MD, USA.

29. **Practical and Accurate Runtime Application Protection against DoS Attacks**
Mohamed Elsabagh, Dan Fleck, Angelos Stavrou, Michael Kaplan, Thomas Bowen
In the Proceedings of 20th International Symposium on Research on Attacks, Intrusions
and Defenses (RAID 2017). September 18-20, 2017, Atlanta, Georgia, USA.

30. **E-Android: A New Energy Profiling Tool for Smartphones**
Xing Gao, Dachuan Liu, Daiping Liu, Haining Wang, Angelos Stavrou

In the proceedings of the the 37th IEEE International Conference on Distributed Computing Systems (ICDCS 2017), June 5-8, 2017, Atlanta, Georgia, USA. (Acceptance Rate: 16.9%, 90/531)

31. **Detecting Passive Cheats in Online Games via Performance-Skillfulness Inconsistency**
Daiping Liu, Xing Gao, Mingwei Zhang, Haining Wang, Angelos Stavrou
In the proceedings of the 47th IEEE/IFIP International Conference on Dependable Systems
and Networks (DSN 2017), June 26-29, 2017, Denver, Colorado, USA.
(Acceptance Rate: 22.3%, 49/220)

32. **Strict Virtual Call Integrity Checking for C++ Binaries** (**Distinguished paper award**)
Mohamed Elsabagh, Dan Fleck, Angelos Stavrou
In the Proceedings of the ACM Asia Conference on Computer and Communications Security
(ASIACCS) 2017, April 2-6, 2017 Abu Dhabi, UAE (Acceptance Rate: 18.7%, 67/359)

33. **Why Software DoS is Hard to Fix: Denying Access in Embedded Android Platforms**
Ryan Johnson, Mohamed Elsabagh, and Angelos Stavrou
In the proceedings of the 14th International Conference on Applied Cryptography and Network Security (ACNS) 2016, June 19-22, 2016, London, UK (Acceptance Rate: 19.13%, 35/183).

34. **When a Tree Falls: Using Diversity in Ensemble Classifiers to Identify Evasion in Malware Detectors**
Charles Smutz and Angelos Stavrou
In the proceedings of the Network and Distributed System Security Symposium (NDSS) 2016, February 21-24, San Diego, California, USA (Acceptance Rate: 15.4%, 60/389).

35. **Targeted DoS on Android: How to Disable Android in 10 Seconds or Less**
Ryan Johnson, Mohamed Elsabagh, Angelos Stavrou, and Vincent Sritapan
In the proceedings of the 10th Malware Conference (MALCON) Oct. 2015, IEEE Computer Society
ISBN: 978-1-5090-0317-4 pp: 136-143 Puerto Rico, USA.

36. **Preventing Exploits in Microsoft Office Documents through Content Randomization**
Charles Smutz and Angelos Stavrou
In the proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses (RAID), November 2015, Kyoto, Japan (Acceptance Rate:

23.5%, 28/119).

37. **Continuous Authentication on Mobile Devices Using Power Consumption, Touch Gestures and Physical Movement of Users**
Rahul Murmuria, Angelos Stavrou, Daniel Barbara, Dan Fleck
In the proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses (RAID), November 2015, Kyoto, Japan (Acceptance Rate: 23.5%, 28/119).

38. **Privacy Risk Assessment on Online Photos**
Haitao Xu, Haining Wang, Angelos Stavrou
In the proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses (RAID), November 2015, Kyoto, Japan (Acceptance Rate: 23.5%, 28/119).

39. **Radmin: Early Detection of Application-Level Resource Exhaustion and Starvation Attacks**
Mohamed Elsabagh, Daniel Barbara, Daniel Fleck, Angelos Stavrou
In the proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses (RAID), November 2015, Kyoto, Japan (Acceptance Rate: 23.5%, 28/119).

40. **On the DNS Deployment of Modern Web Services (Best paper nominee)**
Shuai Hao, Haining Wang, Angelos Stavrou, and Evgenia Smirni
In the proceeding of the 23rd IEEE International Conference on Network Protocols (ICNP)
November 10-13 2015, San Francisco, CA, USA (Acceptance rate: 20%).

41. **Analysis of Content Copyright Infringement in Mobile Application Markets (Best paper award)**
Ryan Johnson, Nikolaos Kiourtis, Angelos Stavrou, and Vincent Sritapan
In the proceedings of APWG/IEEE eCrime Research Summit 2015, May 2015, Barcelona, Spain.

42. **Using Hardware Features for Increased Debugging Transparency**
Fengwei Zhang, Kevin Leach, Angelos Stavrou, Haining Wang, and Kun Sun
In the Proceedings of the 36th IEEE Symposium on Security and Privacy (Oakland 2015), May 2015, San Jose, CA (Acceptance Rate: 13.5%, 55/407).

43. **Resurrecting the READ_LOGS Permission on Samsung Devices**
Ryan Johnson and Angelos Stavrou
In the briefings of Blackhat Asia 2015.

44. **E-commerce Reputation Manipulation: The Emergence of Reputation-Escalation-as-a-Service (Best paper nominee)**
Haitao Xu, Daiping Liu, Haining Wang and Angelos Stavrou

In the Proceedings of 24th World Wide Web Conference (WWW 2015) (Acceptance Rate: 14.1%, 131/929).

45. **TrustLogin: Securing Password-Login on Commodity Operating Systems**
Fengwei Zhang, Kevin Leach, Haining Wang, and Angelos Stavrou
In the Proceedings of The 10th ACM Symposium on Information, Computer and Communications Security (AsiaCCS'15), Singapore, April 2015 (Acceptance Rate: 17.8%, 48/269).

46. **transAD: An Anomaly Detection Network Intrusion Sensor for the Web (short paper)**
Sharath Hiremagalore, Daniel Barbara, Dan Fleck, Walter Powell, and Angelos Stavrou
In the Proceedings of Information Security Conference (ISC 2014), Lecture Notes in Computer Science p 477-489, Hong Kong, Oct 2014. (Acceptance Rate: 17.8%, 48/269)

47. **A Framework to Secure Peripherals at Runtime**
Fengwei Zhang, Haining Wang, Kevin Leach, Angelos Stavrou
European Symposium on Research in Computer Security (ESORICS) p. 219-238 (2014)
(Acceptance Rate: 24.8%, 58/234)

48. **Click Fraud Detection on the Advertiser Side**
Haitao Xu, Daiping Liu, Aaron Koehl, Haining Wang, Angelos Stavrou
European Symposium on Research in Computer Security (ESORICS) p. 419-438 (2014)
(Acceptance Rate: 24.8%, 58/234)

49. **Activity Spoofing and Its Defense in Android Smartphones**
Brett Cooley, Haining Wang, and Angelos Stavrou
In the proceedings of the 12th International Conference on Applied Cryptography and Network Security (*ACNS 2014*) Lausanne, Switzerland. (Acceptance Rate: 22.5%, 33/147)

50. **Catch Me if You Can: A Cloud-Enabled DDoS Defense**
Quan Jia, Huangxin Wang, Dan Fleck, Fei Li, Angelos Stavrou, Walter A. Powell.
In the Proceedings of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (IEEE DSN 2014), Atlanta, Georgia USA, June 23 - 26, 2014.

51. **Detecting Malicious Javascript in PDF through Document Instrumentation**
Daiping Liu, Haining Wang, and Angelos Stavrou.
In the Proceedings of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (IEEE DSN 2014), Atlanta, Georgia USA,

June 23 - 26, 2014.

52. **PyTrigger: A System to Trigger & Extract User-Activated Malware Behavior**
Dan Fleck, Arnur Tokhtabayev, Alex Alarif, Angelos Stavrou, and Tomas Nykodym.
In the proceedings of the 8th ARES Conference (ARES 2013),
University of Regensburg, Germany September 2nd - 6th, 2013.

53. **MOTAG: Moving Target Defense Against Internet Denial of Service Attacks**
Quan Jia, Kun Sun, Angelos Stavrou.
In the proceedings of the International Conference on Computer Communications and Networks ICCCN 2013 Nassau, Bahamas July 30 - August 2, 2013.

54. **Behavioral Analysis of Android Applications Using Automated Instrumentation**
Mohammad Karami, Mohamed Elsabagh, Parnian Najafiborazjani, and Angelos Stavrou.
In the Proceedings of the 7th International Conference on Software Security and Reliability
(IEEE SERE 2013), 18-20 June 2013, Washington DC, USA. (Acceptance rate 30%)

55. **Forced-Path Execution for Android Applications on x86 Platforms**
Ryan Johnson, and Angelos Stavrou.
In the Proceedings of the 7th International Conference on Software Security and Reliability
(IEEE SERE 2013), 18-20 June 2013, Washington DC, USA. (Acceptance rate 30%)

56. **Towards a Cyber Conflict Taxonomy**
Scott Applegate and Angelos Stavrou.
To appear in the Proceedings of the 5th International Conference on Cyber Conflict (CyCon 2013)
NATO Cooperative Cyber Defence Centre of Excellence conference, 4-7 June 2013 in Tallinn, Estonia.

57. **Spectre: A Dependable Introspection Framework via System Management Mode**
Fengwei Zhang, Kevin Leach, Kun Sun, and Angelos Stavrou.
In the Proceedings of the 43nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (IEEE DSN 2013), Budapest, 24 - 27 June 2013.
(Acceptance Rate: 19.6%)

58. **Exposing Software Security and Availability Risks For Commercial Mobile Devices (CMDs)**
Ryan Johnson, Zhaohui Wang, Angelos Stavrou, and Jeff Voas.
In the Proceedings of the IEEE RAMS 2013, Orlando, Florida, 28 - 31 January 2013.

59. **Malicious PDF Detection Using Metadata and Structural Features**
Charles Smutz and Angelos Stavrou.
In the Proceedings of the 2012 Annual Computer Security Applications Conference

(ACSAC), Orlando, Florida, USA, December 3-7, 2012. (Acceptance Rate: 19%, 44/231)

60. **Malware Characterization using Behavioral Components**
Chaitanya Yavvari, Arnur Tokhtabayev, Huzefa Rangwala, and Angelos Stavrou.
In the Proceedings of 6th International Conference "Mathematical Methods, Models, and Architectures for Computer Network Security", St. Petersburg, Russia, October 17-20, 2012.

61. **Exposing Security Risks for Commercial Mobile Devices**(Invited)
Zhaohui Wang, Ryan Johnson, Rahul Murmuria, and Angelos Stavrou.
In the Proceedings of 6th International Conference "Mathematical Methods, Models, and Architectures for Computer Network Security", St. Petersburg, Russia, October 17-20, 2012.

62. **Mobile Application and Device Power Usage Measurements**
Rahul Murmuria, Jeffrey Medsger, Angelos Stavrou.
In the Proceedings of the 6th International Conference on Software Security and Reliability (SERE 2012), Washington, DC, June 2012.

63. **Netgator: Malware Detection Using Program Interactive Challenges**
Brian Schulte, Haris Andrianakis, Kun Sun, Angelos Stavrou.
In the Proceedings of the 9th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2012), Heraklion, Crete, Greece, July 26-27th, 2012.

64. **A Dependability Analysis of Hardware-Assisted Polling Integrity Checking Systems**
Jiang Wang, Kun Sun, and Angelos Stavrou.
In the Proceedings of the 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2012), Boston, Massachusetts, June, 2012.

65. **Implementing & Optimizing an Encryption File System on Android**
Zhaohui Wang, Rahul Murmuria, and Angelos Stavrou.
In the Proceedings of the IEEE International Conference on Mobile Data Management
(IEEE MDM 2012), July 23 - 26, 2012, Bangalore, India. (Acceptance Rate: 22/88)

66. **Analysis Android Applications' Permissions** (short paper)
Ryan Johnson, Zhaohui Wang, Corey Gagnon and Angelos Stavrou.
In the Proceedings of the 6th International Conference on Software Security and Reliability (SERE 2012), Washington, DC, June 2012.

67. **Mutual Authentication for USB Communications** (short paper)
Zhaohui Wang, Ryan Johnson and Angelos Stavrou.
In the Proceedings of the 6th International Conference on Software Security and Reliability (SERE 2012), Washington, DC, June 2012.

68. **A Framework for Automated Security Testing of Android Applications on the Cloud** (short)
Sam Malek, Naeem Esfahani, Thabet Kacem, Riyadh Mahmood, Nariman Mirzaei, and Angelos Stavrou. In the Proceedings of the 6th International Conference on Software Security and Reliability (SERE 2012), Washington, DC, June 2012.

69. **SecureSwitch: BIOS-Assisted Isolation and Switch between Trusted and Untrusted Commodity OSes**
Kun Sun, Jiang Wang, Fengwei Zhang and Angelos Stavrou.
In the Proceedings of the 19th Annual Network & Distributed System Security Symposium NDSS 2012, San Diego, California, 5-8 February 2012. Impact Factor: 2.60 (Acceptance Rate: 46/258 - 17.8%). [Presentation]

70. **Hardware-Assisted Application Integrity Monitor**
Jiang Wang, Kun Sun, Angelos Stavrou.
In the Proceedings of IEEE Hawaii International Conference on System Sciences (HICSS45) pp. 5375-5383, 45th Hawaii International Conference on System Sciences, 2012 January 4-7, 2012, Grand Wailea, Maui, USA. Impact Factor: N/A, (Acceptance Rate: N/A).

71. **Cross-domain Collaborative Anomaly Detection: So Far Yet So Close**
Nathaniel Boggs, Sharath Hiremagalore, Angelos Stavrou, Salvatore J. Stolfo.
In the Proceedings of the 14th International Symposium on Recent Advances in Intrusion Detection (RAID). September 2011, Menlo Park, CA.
Impact Factor: 2.20 (Acceptance rate: 20/87 - 23%).

72. **Trading Elephants For Ants: Efficient Post-Attack Reconstitution (Short paper)**
Meixing Le, Zhaohui Wang, Quan Jia, Angelos Stavrou, Anup Ghosh and Sushil Jajodia
In the Proceedings of the 7th International ICST Conference on Security and Privacy in Communication Networks (Securecomm 2011), p. 1-10, September 7-9 2011, London.
Impact Factor: N/A, (Acceptance rate: 24%).

73. **Predicting Network Response Times Using Social Information (short paper)**
Chen Liang, Sharath Hiremagalore, Angelos Stavrou and Huzefa Rangwala.
In the Proceedings of the ACM 2011 Conference on Advances in Social Networks Analysis and Mining, p. 527-531, July, 2011, Kaohsiung, Taiwan. Impact Factor: N/A, (Acceptance rate: 25%) [bib]

74. **Breaching and Protecting an Anonymizing Network System**
Jason Clark and Angelos Stavrou.
In the Proceedings of the 6th Annual Symposium on Information Assurance (ASIA 2011).
Impact Factor: N/A, (Acceptance rate: 50%).

75. **Advantages and disadvantages of remote asynchronous usability testing using amazon mechanical turk**
Erik Nelson and Angelos Stavrou.
Proceedings of the Human Factors and Ergonomics Society 55th Annual Meeting, pages 1080-1084, HFES 2011 Conference, Red Rock Resort, Las Vegas, Nevada, September 19-23, 2011.
Impact Factor: N/A, (Acceptance rate: N/A).

76. **Exploiting Smart-Phone USB Connectivity For Fun And Profit (Extended Version)**
Angelos Stavrou and Zhaohui Wang.
BlackHat Technical Conference DC 2011 - Technical Briefings Session.

77. **Exploiting Smart-Phone USB Connectivity For Fun And Profit**
Zhaohui Wang and Angelos Stavrou.
In the Proceedings of the 26th Annual Computer Security Applications Conference (ACM ACSAC)
p. 357-366. December 6-10, 2010, Austin, Texas, USA. Impact Factor: 1.82 (Acceptance rate: 39/227) [bib]

78. **Experimental Results of Cross-Site Exchange of Web Content Anomaly Detector Alerts**
Nathaniel Boggs, Sharath Hiremagalore, Angelos Stavrou, and Salvatore J. Stolfo.
In the Proceedings of IEEE Conference on Homeland Security Technologies (IEEE HST 2010),
November 8-10, 2010, Waltham, MA, USA. Impact Factor: N/A (Acceptance rate: N/A).

79. **An Adversarial Evaluation of Network Signaling and Control Mechanisms**
Kangkook Jee, Stelios Sidiroglou-Douskos, Angelos Stavrou, and Angelos D. Keromytis.
In the Proceedings of the 13th International Conference on Information Security and Cryptology (ICISC).
December 2010, Seoul, Korea. Impact Factor: N/A (Acceptance rate: N/A).

80. **Small World VoIP**
Xiaohui Yang, Angelos Stavrou, Ram Dantu, and Duminda Wijesekera.
In the Proceedings of the Second International Conference on Mobile Computing, Applications,
and Services MobiCASE, October 25-28, 2010, Santa Clara, CA, USA.
Impact Factor: N/A (Acceptance rate: N/A).

81. **QoP and QoS policy cognizant policy composition**
Paul Seymer, Angelos Stavrou, Duminda Wijesekera, Sushil Jajodia.
In the Proceedings of the IEEE International Symposium on Policies for Distributed Systems

and Networks, p. 77-86, Fairfax, VA, July 21-23, 2010. (Acceptance rate: 19.2%) [bib]

82. **Providing Mobile Users' Anonymity in Hybrid Networks**
Claudio Ardagna, Sushil Jajodia, Pierangela Samarati, and Angelos Stavrou (Alphabetic)
In the Proceedings of the 15th European Symposium on Research in Computer Security
(ESORICS 2010), p. 540-557, September 2010, Athens, Greece.
Impact Factor: 1.45 (Acceptance rate: 42/210 - 20%). [bib]

83. **Traffic Analysis Against Low-Latency Anonymity Networks Using Available Bandwidth Estimation**
Sambuddho Chakravarty, Angelos Stavrou, and Angelos D. Keromytis.
In the Proceedings of the 15th European Symposium on Research in Computer Security
(ESORICS 2010) p. 249-267, September 2010, Athens, Greece.
Impact Factor: 1.45 (Acceptance rate: 42/210 - 20%). [bib]

84. **HyperCheck: A Hardware-Assisted Integrity Monitor**
Jiang Wang, Angelos Stavrou, and Anup K. Ghosh.
In the Proceedings of 13th International Symposium on Recent Advances in Intrusion Detection
(RAID 2010), p. 158-177, Ottawa, Canada, September 15-17, 2010.
Impact Factor: 2.20 (Acceptance rate: 24/104 - 23.1%). [bib]

85. **A Virtualization Architecture for In-Depth Kernel Isolation**
Jiang Wang, Sameer Niphadkar, Angelos Stavrou, Anup K. Ghosh.
In the Proceedings of 43rd Hawaii International International Conference on Systems Science,
IEEE Computer Society, p. 1-10, 5-8 January 2010, Koloa, Kauai, HI, USA.
Impact Factor: N/A (Acceptance rate: N/A).

86. **Privacy preservation over untrusted mobile networks**
Claudio A. Ardagna, Sushil Jajodia, Pierangela Samarati, Angelos Stavrou in Privacy in Location-Based
Applications: Research Issues and Emerging Trends, Springer Lecture Notes in Computer Science, Volume
5599, 2009, pages 84-105. Impact Factor: N/A (Acceptance rate: N/A).

87. **Deny-by-Default Distributed Security Policy Enforcement in Mobile Ad Hoc Networks (short)**
Mansoor Alicherry, Angelos D. Keromytis, and Angelos Stavrou.
In the Proceedings of the 5th International ICST Conference on Security and Privacy in Communication
Networks SECURECOMM 2009, p. 41-50. September 2009, Athens, Greece.
Impact Factor: N/A, (Acceptance rate: 25.3%). [bib]

88. **Adding Trust to P2P Distribution of Paid Content**
Alex Sherman, Angelos Stavrou, Jason Nieh, Angelos D. Keromytis, and Clifford Stein.
In the Proceedings of the 12th Information Security Conference (ISC), p.459-474. September 2009, Pisa, Italy. Impact Factor: 1.24, (Acceptance rate: 27.6%). [bib]

89. **A2M: Access-Assured Mobile Desktop Computing**
Angelos Stavrou, Ricardo A. Baratto, Angelos D. Keromytis, and Jason Nieh.
In the Proceedings of the 12th Information Security Conference (ISC), p. 186-201. September 2009, Pisa, Italy. Impact Factor: 1.24, (Acceptance rate: 27.6%). [bib]

90. **Adaptive Anomaly Detection via Self-Calibration and Dynamic Updating**
Gabriela F. Cretu, Angelos Stavrou, Michael E. Locasto, Salvatore J. Stolfo.
In the Proceedings of 12th International Symposium On Recent Advances In Intrusion Detection,
p. 41-60. Saint-Malo, Brittany, France, September 23-25, 2009.
Impact Factor: 2.20 (Acceptance rate: 17 / 59 - 28.8%). [bib]

91. **SQLProb: A Proxy-based Architecture towards Preventing SQL Injection Attacks**
Anyi Liu, Yi Yuan, Duminda Wijesekera, and Angelos Stavrou.
In the Proceedings of 24th Annual ACM Symposium on Applied Computing (SAC'09), p. 2054-2061
March 8-12, 2009, Honolulu, Hawaii. Impact Factor: N/A, (Acceptance Rate: 16.6%). [bib]

92. **A Security Architecture for Information Assurance and Availability in MANETs**
Angelos Stavrou, and Anup K. Ghosh.
In the Proceedings of IEEE Conference on Military Communications (MILCOM '08), p. 1 - 8, November 2008, San Diego, CA. Impact Factor: N/A, (Acceptance Rate: N/A).
Impact Factor: N/A, (Acceptance Rate: N/A). [bib]

93. **PAR: Payment for Anonymous Routing**
Elli Androulaki, Mariana Raykova, Shreyas Srivatsan, Angelos Stavrou, and Steven M. Bellovin.
In the Proceedings of 8th Privacy Enhancing Technologies Symposium, p. 219-236, Leuven, Belgium
July 23 - July 25, 2008. Impact Factor: 1.95, Acceptance rate: 13/49 - 26%). [bib]

94. **The Hidden Difficulties of Watching and Rebuilding Networks.**
Michael Locasto and Angelos Stavrou.
IEEE Security and Privacy, vol. 6, no. 2, pp. 79-82, Mar/Apr, 2008.
Impact Factor: 1.17, (Acceptance Rate: N/A). [bib]

95. **Pushback for Overlay Networks: Protecting against Malicious Insiders**
Angelos Stavrou, Michael E. Locasto, and Angelos D. Keromytis.

In the Proceedings of the 6th International Conference on Applied Cryptography and Network Security (ACNS). June 2008, New York, NY.
Impact Factor: N/A, (Acceptance Rate: N/A). [bib]

96. **Casting out Demons: Sanitizing Training Data for Anomaly Sensors**
Gabriela F. Cretu, Angelos Stavrou , Michael E. Locasto, Salvatore J. Stolfo, and Angelos D. Keromytis.
In the Proceedings of the IEEE Symposium on Security & Privacy p. 81-95. May 2008, Oakland, CA.
Impact Factor: 4.15, (Acceptance Rate: 11.2%) [bib]

97. **On the Infeasibility of Modeling Polymorphic Shellcode**
Yingbo Song, Michael E. Locasto, Angelos Stavrou , Angelos D. Keromytis, and Salvatore J. Stolfo.
In the Proceedings of the 13th ACM Conference on Computer and Communications Security (CCS),
pp. 541 - 551. October/November 2007, Alexandria, VA. Impact Factor: 2.87, (Acceptance rate: 18.1%) [bib]

98. **A Study of Malcode-Bearing Documents**
Weijen Li, Salvatore Stolfo, Angelos Stavrou, Elli Androulaki, and Angelos D. Keromytis.
In Proceedings of the 4th GI International Conference on Detection of Intrusions & Malware, and
Vulnerability Assessment (DIMVA), pp. 231 - 250. July 2007, Lucerne, Switzerland.
Impact Factor: 1.42, (Acceptance rate: 21%). [bib]

99. **From STEM to SEAD: Speculative Execution for Automated Defense.**
Michael E. Locasto, Angelos Stavrou, Gabriela F. Cretu, and Angelos D. Keromytis.
In the Proceedings of the USENIX Annual Technical Conference (USENIX 2007), pp. 219-232, June 2007, Santa Clara, CA. Impact Factor: N/A, (Acceptance rate: 18.75%) [bib]

100. **Network Security as a Composable Service**
Stelios Sidiroglou, Angelos Stavrou, and Angelos D. Keromytis.
In the Proceedings of the IEEE Sarnoff Symposium. January 2007, Princeton, NJ. (Invited paper)

101. **Countering DDoS Attacks with Multi-path Overlay Networks**
Angelos Stavrou and Angelos D. Keromytis.
In the Information Assurance Technology Analysis Center (IATAC) Information Assurance
Newsletter (IAnewsletter), vol. 9, no. 3, November 2006. (Invited paper, based on the CCS 2005 paper.)[pdf]

102. **W3Bcrypt: Encryption as a Stylesheet**
Angelos Stavrou, Michael E. Locasto, and Angelos D. Keromytis. In the Proceedings

of the
4th International Conference on Applied Cryptography and Network Security (ACNS 2006),
pp. 349-364, June 6-9, 2006, Singapore. Impact Factor: 1.44, (Acceptance rate: 33 / 218 - 15.1%) [bib]

103. **Countering DoS Attacks With Stateless Multipath Overlays**
Angelos Stavrou and Angelos D. Keromytis.
In the Proceedings of the 12th ACM Conference on Computer and Communications Security (CCS),
pp. 249 - 259. November 2005, Alexandria, VA. Impact Factor: 2.87, (Acceptance rate: 15.2%) [bib]

104. **gore: Routing-Assisted Defense Against DDoS Attacks**
Stephen T. Chou, Angelos Stavrou, John Ioannidis, and Angelos D. Keromytis.
In the Proceedings of the 8th Information Security Conference (ISC), p. 179-193. September 2005, Singapore. Impact Factor: 1.24, (Acceptance rate: 14%). [bib]

105. **MOVE: An End-to-End Solution To Network Denial of Service**
Angelos Stavrou, Angelos D. Keromytis, Jason Nieh, Vishal Misra, and Dan Rubenstein.
In the Proceedings of the Internet Society (ISOC) Symposium on Network and Distributed Systems Security
(NDSS), pp. 81 - 96. February 2005, San Diego, CA. Impact Factor: 2.60
(Acceptance rate: 12.9%). [bib]

106. **Content distribution for seamless transmission**
Edward G. Coffman Jr., Andreas Constantinides, Dan Rubenstein, Bruce Shepherd, Angelos Stavrou
In the Proceedings of SIGMETRICS Performance Evaluation Review 32(2): 31-32 (2004) [pdf (936 KB)].

107. **A Pay-per-Use DoS Protection Mechanism For The Web**
Angelos Stavrou, John Ioannidis, Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein.
In the Proceedings of the Applied Cryptography and Network Security (ACNS) Conference.
June 2004, Yellow Mountain, China. LNCS Volume 3089/2004, pp. 120-134, ISBN: 3-540-22217-0.
Impact Factor: 1.44 (Acceptance rate: 12%). [bib]

108. **Using Graphic Turing Tests to Counter Automated DDoS Attacks Against Web Servers**
William G. Morein, Angelos Stavrou, Debra L. Cook, Angelos D. Keromytis, Vishal Misra, Dan Rubenstein.
In the Proceedings of the 10th ACM International Conference on Computer and Communications Security

(CCS), Washington, DC, October 2003. Impact Factor: 2.87, (Acceptance rate: 13.8%) [bib]

109. **A Lightweight, Robust P2P System to Handle Flash Crowds**
Angelos Stavrou, Dan Rubenstein and Sambit Sahu.
In the Proceedings of IEEE ICNP 2002, Paris, France, November, 2002.
[Proceedings Version ps (252K)] [Proceedings Version ps.gz (65K)] [Proceedings Version pdf (143K)]
An earlier version is available as Columbia Technical Report EE020321-1, February, 2002.
[Tech Report ps (508K)] [Tech Report ps.gz (109K)] [Tech Report pdf (242K)].
Impact Factor: N/A, (Acceptance rate: 14.7%). [bib]

Books/Book Chapters

1. **Overlay-Based DoS Defenses**
Angelos Stavrou. In Henk C.A. van Tilborg and Sushil Jajodia, editors, Encyclopedia of Cryptography and Security, 2nd Edition. Springer, 2010.

2. **TCP Modulation Attacks**
Angelos Stavrou. In Henk C.A. van Tilborg and Sushil Jajodia, editors, Encyclopedia of Cryptography and Security, 2nd Edition. Springer, 2010.

Workshops

1. **Microservices made attack-resilient using unsupervised service fissioning [PDF]**
Ataollah Fatahi Baarzi, George Kesidis, Dan Fleck, Angelos Stavrou.
In EuroSec '20: Proceedings of the 13th European workshop on Systems Security, April 2020 Pages 31–36.

2. **Breaking BLE Beacons For Fun But Mostly Profit**
Constantinos Kolias, Lucas Copi, Fengwei Zhang, Angelos Stavrou
EUROSEC 2017: 4:1-4:6

3. **Your Data in Your Hands: Privacy-preserving User Behavior Models for Context Computation**
Rahul Murmuria, Angelos Stavrou, Daniel Barbara, and Vincent Sritapan
To appear in the proceedings of International Workshop on Behavioral Implications of Contextual Analytics (co-located with IEEE PerCom 2017)

4. **Authentication Feature and Model Selection using Penalty Algorithms**
Rahul Murmuria and Angelos Stavrou.
In the proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS 2016). USENIX Association, 2016, Way workshop, Denver Colorado June 22-24, 2016

5. **Switchwall: Automated Topology Fingerprinting & Behavior Deviation Identification**
   Nelson Nazzicari, Javier Almillategui, Angelos Stavrou and Sushil Jajodia.
   In the Proceedings of the 8th International Workshop on Security and Trust Management (STM 2012)
   in conjunction with ESORICS 2012, Pisa, Italy - September 13-14, 2012

6. A Whitebox Approach for Automated Security Testing of Android Applications on the Cloud
   Riyadh Mahmood, Naeem Esfahani, Thabet Kacem, Nariman Mirzaei, Sam Malek, and Angelos Stavrou.
   In the Proceedings of the 7th International Workshop on Automation of Software Test (AST 2012), Zurich, Switzerland, June 2012.

7. The MEERKATS Cloud Security Architecture
   Angelos D. Keromytis, Roxana Geambasu, Simha Sethumadhavan, Salvatore J. Stolfo, Junfeng Yang, Azzedine Benameur, Marc Dacier, Matthew Elder, Darrell Kienzle, and Angelos Stavrou.
   In the Proceedings of the 3rd International Workshop on Security and Privacy in Cloud Computing
   (ICDCS-SPCC). June 2012, Macao, China.

8. CapMan: Capability-based Defense against Multi-Path Denial of Service (DoS) Attacks in MANET
   Quan Jia, Kun Sun and Angelos Stavrou.
   In the Proceedings of the First International Workshop on Privacy, Security and Trust in Mobile and Wireless Systems
   (MobiPST 2011) in conjunction with *20th International Conference on Computer Communications and Networks (ICCCN 2011)*

9. The MINESTRONE Architecture: Combining Static and Dynamic Analysis Techniques for Software Security
   Angelos D. Keromytis, Salvatore J. Stolfo, Junfeng Yang, Angelos Stavrou, Anup Ghosh, Dawson Engler, Marc Dacier, Matthew Elder, and Darrell Kienzle.
   In the Proceedings of the 1st Workshop on Systems Security (SysSec).
   July 2011, Amsterdam, Netherlands.

10. Firmware-assisted Memory Acquisition and Analysis tools for Digital Forensic (short paper)
    Jiang Wang, Fengwei Zhang, Kun Sun, and Angelos Stavrou.
    In the Sixth International Workshop on Systematic Approaches to Digital Forensic Engineering (IEEE SADFE 2011).
    In conjunction with IEEE Security and Privacy Symposium, Oakland, CA, USA, May 26, 2011

11. Moving Forward, Building An Ethics Community (Panel Statements) - Computer Security Ethics, Quo Vadis?

Erin Kenneally, Angelos Stavrou, John McHugh, and Nicolas Christin.
In the proceedings of the 2nd Workshop on Ethics in Computer Security Research 2011
Springer Lecture Notes in Computer Science (LNCS).

12. Scalable Web Object Inspection and Malfease Collection
Charalampos Andrianakis, Paul Seymer, and Angelos Stavrou.
In the Proceedings of the 5th USENIX Workshop on Hot Topics in Security (HotSec '10).
August 10, 2010 Washington, DC. (Acceptance rate: 11/57)

13. Fine-grained Sharing of Health Records using XSPA Profile for XACML
A. Al-Faresi, Bo Yu, Khalid Moidu, Angelos Stavrou, Duminda Wijesekera, Anoop Singhal
In the Proceedings of 1st USENIX Workshop on Health Security and Privacy (HealthSec '10),
August, 2010, Washington DC, USA.

14. Evaluating a Collaborative Defense Architecture for MANETs
Mansoor Alicherry, Angelos Stavrou, and Angelos D. Keromytis.
In the Proceedings (electronic) of the IEEE Workshop on Collaborative Security Technologies (CoSec),
pp. 37 - 42. December 2009, Bangalore, India. (Acceptance rate: 17.2%).

15. Keep your friends close: the necessity for updating an anomaly sensor with legitimate environment changes.
Angelos Stavrou, Gabriela F. Cretu, Michael E. Locasto, Salvatore J. Stolfo.
In the Proceedings of the 2nd ACM Workshop on Security and Artificial intelligence (Chicago, Illinois, USA, November 09 - 09, 2009). AISec '09. ACM, New York, NY, 39-46. (Position paper)

16. The Heisenberg Measuring Uncertainty in Lightweight Virtualization Testbeds
Quan Jia, Zhaohui Wang and Angelos Stavrou.
In the Proceedings of 2nd Workshop on Cyber Security Experimentation and Test (CSET '09). August, 2009, Montreal, Canada.

17. Universal Multi-Factor Authentication Using Graphical Passwords
Alireza Pirayesh Sabzevar, and Angelos Stavrou.
In the Proceedings of the 2nd IEEE/ACM Workshop on Security and Privacy in Telecommunications and Information Systems (SePTIS). December 2008, Bali, Indonesia.

18. Identifying Proxy Nodes in a Tor Anonymization Circuit
Sambuddho Chakravarty, Angelos Stavrou, and Angelos D. Keromytis.
In the Proceedings of the 2nd IEEE/ACM Workshop on Security and Privacy in Telecommunications and Information Systems (SePTIS). December 2008, Bali, Indonesia.

19. A multi-path approach for k-anonymity in mobile hybrid networks
Claudio Agostino Ardagna, Angelos Stavrou, Sushil Jajodia, Pierangela Samarati and Rhandi Martin.
In the Proceedings of International Workshop on Privacy in Location-Based Applications (PiLBA '08), October 2008.

20. Efficiently Tracking Application Interactions using Lightweight Virtualization
Yih Huang, Angelos Stavrou, Anup K. Ghosh and Sushil Jajodia.
In the Proceeding of the 1st Workshop on Virtualization Security (VMSec), in conjunction
with ACM CCS 2008, October 2008.

21. Return Value Predictability for Self-Healing
Michael E. Locasto, Angelos Stavrou>, Gabriela F. Cretu, Angelos D. Keromytis, and Salvatore J. Stolfo.
In the Proceedings of the 3rd International Workshop on Security (IWSEC), November 2008, Kagawa, Japan.

22. Online Training and Sanitization of AD Systems (extended abstract)
Gabriela F. Cretu, Angelos Stavrou, Michael E. Locasto, Salvatore J. Stolfo.
In the Proceedings of NIPS 2007 Workshop on Machine Learning in Adversarial Environments for Computer Security,
December 2007, Vancouver, B.C., Canada. [pdf]

23. Data Sanitization: Improving the Forensic Utility of Anomaly Detection Systems
Gabriela F. Cretu, Angelos Stavrou, Salvatore J. Stolfo, Angelos D. Keromytis.
In the Proceedings of the 3rd Workshop on Hot Topics in System Dependability (HotDep), pp. 64 - 70. June 2007, Edinburgh, UK. [pdf]

24. Bridging the Network Reservation Gap Using Overlays
Angelos Stavrou, David Turner, Angelos D. Keromytis, and Vassilis Prevelakis.
In the Proceedings of the 1st Workshop on Information Assurance for Middleware Communications (IAMCOM).
January 2007, Bangalore, India. [pdf] [ps]

25. Dark Application Communities
Michael E. Locasto, Angelos Stavrou, and Angelos D. Keromytis.
In the Proceedings of the 15th New Security Paradigms Workshop (NSPW 2006).
September 2006, Schloss Dagstuhl, Germany. [pdf] [ps]

Technical Reports

1. Netgator: Malware Detection Through Program Interactive Proofs
Brian Schulte, Rhandi Martin, Haris Andrianakis and Angelos Stavrou, GMU-CS-TR-2011-6

2. SecureSwitch: BIOS-Assisted Isolation and Switch between Trusted and Untrusted Commodity OSes
   Kun Sun, Jiang Wang, Fengwei Zhang and Angelos Stavrou, GMU-CS-TR-2011-7

3. An Analysis of System Management Mode (SMM)-based Integrity Checking Systems and Evasion Attacks
   Jiang Wang, Kun Sun and Angelos Stavrou, GMU-CS-TR-2011-8

4. Ruminate: A Scalable Architecture for Deep Network Analysis
   Charles Smutz and Angelos Stavrou, GMU-CS-TR-2010-20.

Professional Activities & Service

Founder & CEO

- **Kryptowire LLC (https://www.kryptowire.com/)**

Founder & Chief Scientist

- **Quokka INC (https://www.quokka.io)**

Editorial Positions, Panels, and Boards

- IEEE Security and Privacy Magazine (01-01-2017 - Present)
  Type: Internet publication
  Role: Associate or guest editor/curator
- International Journal of Information Security (01-01-2019 - Present)
  Type: Journal
  Role: Editorial/curatorial board member
- IEEE Transactions on Computers (01-01-2020 - Present)
  Type: Journal
  Role: Associate or guest editor/curator
  Description: Associate Editor of IEEE Transactions on Computers
- IEEE Internet Computing (01-01-2020 - Present)
  Type: Journal
  Role: Associate or guest editor/curator

Past:
Associate Editor, IEEE Transactions on Reliability, September 2015 - 2020
IET Journal on Information Security, May 2010 – May 2018
Encyclopedia of Cryptography and Security, Editorial Board Member, March 2010 – 2020

Program Organization:

Program General co-Chair, ACM Conference on Computer and Communications Security (CCS) 2022

Program co-Chair, 10th European Workshop on Systems Security (EuroSec): 2017, 2018

Program co-Chair, Research in Attacks, Intrusions and Defenses (RAID) Symposium, RAID 2013, 2014

Student Travel Grant Chair, ACM Conference on Computer and Communications Security (CCS), 2009, 2010

Program co-Chair, Workshop on Cyber Security Experimentation and Test (CSET): 2009, 2010

Program co-Chair, 1st Workshop on Virtual Machine Security (VMSec): 2008, 2009

Program Committee Member (Selected):

USENIX Security Symposium: 2007, 2008, 2009, 2020, 2021
ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec) 2020
IEEE MILCOM (Track 3): 2016, 2017, 2018
IEEE Conference on Dependable and Secure Computing: 2017
IEEE International Conference on Software Quality, Reliability and Security: 2016
Recent Advances in Intrusion Detection (RAID): 2011, 2012, 2017
IEEE Symposium on Security and Privacy (IEEE S&P): 2010, 2011, 2012
ACM Conference on Computer and Communications Security (CCS): 2009, 2010
Annual Computer Security Applications Conference (ACSAC): 2009, 2010, 2011, 2012, 2013
Network and Distributed System Security Symposium (NDSS): 2009, 2010
International Conference on Distributed Computing Systems (ICDCS): 2009, 2010, 2011, 2012, 2013
25th ACM Symposium On Applied Computing (SAC): 2010
Financial Cryptography and Data Security: 2010, 2011, 2012
5th ACM Int'l Conference on emerging Networking EXperiments and Technologies: 2009
USENIX Security Symposium: 2007, 2008, 2009
International ICST Conference on Security and Privacy in Communication Networks (SecureComm), 2009, 2010, 2011
European Workshop on System Security (EUROSEC): 2008, 2009, 2010, 2011
IEEE International Symposium on Policies for Distributed Systems and Networks, POLICY 2010, 2011, 2012
ECML/PKDD Workshop on Privacy and Security issues in Data Mining and Machine Learning PSDML 2010
2nd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET): 2009
European Conference on Computer Network Defense (EC2ND): 2008
Workshop on Cyber Security Experimentation and Test : 2008, 2013
Information Security Conference (ISC): 2008, 2009
European Symposium on Research in Computer Security (ESORICS): 2008
International Workshop on Security and Privacy in Wireless and Mobile Computing, Networking and Communications: 2008

Patent Litigation Experience

Expert witness for patent invalidity reports, depositions, patent office actions, Inter Partes Disputes preparation, product infringement reports, code and system examination.

Law firm: Quinn Emanuel Urquhart & Sullivan, LLP

Fortinet, Inc V. Sophos, Inc

Case No 13-cv-05831-EMC (representing Fortinet)

Deposed for the case.
Project status: Ended

Law firm: Feinberg Day Alberti & Thompson LLP ("Feinberg Day")
Representing: Intellectual Ventures II L.L.C.
Case: Intellectual Ventures II L.L.C. v. JP Morgan Chase & Co. et al (representing IV)
Case: 1:13-cv-03777 | New York Southern District Court
Case: 1:13-cv-03777-AKH
Civil Action No. 2:13-cv-1106-AKK
Case: l:13-CV-02454-WSD
Case: 2:13-CV-04160-NKL
Case: 2:13-CV-785
Case: 8:13-cv-00167
Case: IPR2014-00786, Patent No. 6,826,694
Project status: Ended

IP Patent Consulting

Law firm: Brian Owens, Esq

Case: Multiple pertaining to analysis of patent validity.

Project: On-going consulting on patent infringement, validity, and claim construction

Patent Filling/Examination Experience

Consulted GMU legal team on patent filling, claim structure, and responses to USPTO examination and re-examination of relevant patents.

1. U.S. Patent Application 12/558,841 filed on September 14, 2009, entitled "Distributed Sensor for Detecting Malicious Software."
2. U.S. Patent Application 12/548,175 filed on August 26, 2009, entitled "Event Driven Email Revocation."
3. U.S. Patent Application 12/688,037 filed on January 15, 2010, entitled "Authentication Using Graphical Passwords."
4. U.S. Patent Application 12/757,675 filed on April 9, 2010, entitled "Malware Detector."

5. U.S. Patent Application 12/965,413 filed on December 10, 2010, entitled "Website Detection."
6. U.S. Patent Application 12/835,228 filed on July 13, 2010, entitled "Inferring Packet Management Utility Rules."
7. U.S. Patent Application 611413,673 filed on November 15, 2010, entitled "HyperCheck: A Hardware Assisted Integrity Monitor."
8. U.S. Patent Application 611413,677 filed on November 15, 2010, entitled "Network Traffic Analysis."

Consulted on more patents filled prior to 2007 while I was at Columbia University (worked with Columbia retained law firms).

Advisory Boards, Workshops & Other Professional Activities

Academic Program Director, Masters in Management of Secure Information Systems, George Mason University, 2014 – May 2017
Academic Program Director, Masters in Information Security and Assurance, George Mason University, 2013 - 2015
IEEE Rebooting Computing Committee, 2013 - onwards
Senior Member of the IEEE, 2012 - onwards
Subject Matter Expert, DARPA Transformative Applications, September 2010 - September 2012
USDA Federal Mobile Computing Summit, 2011
NIST Mobile & Smart Phone Technologies Technical Exchange Meeting, 2011
Google Faculty Summit, July 2010
ARO/NSF Workshop on Moving Target Defense, October 2010
National Science Foundation Panels: 2008, 2009
DARPA Cyber Genome Project, Dec 2009
DARPA Digital Object Storage and Retrieval (DOSR), July 2008
DARPA Intrinsically Assurable Mobile Ad-hoc Networks (IAMANETs), January 2008

Ph.D. Thesis Committee Service

1. **Mahmood Riyadh**, Electrical Engineering Department, George Mason University, Summer 2015.
2. **Velegalati Rajesh**, Electrical Engineering Department, George Mason University, Summer 2015.
3. **Yu Bo**, Computer Science Department, George Mason University, Fall 2014.
4. **Jin Jing,** Computer Science Department, George Mason University, Fall 2013.
5. **Xu Min,** Computer Science Department, George Mason University, Fall 2013.
6. **Caixia Wang**, Thesis title: "*Spatial content-based scene matching using a relaxation method*",
   Department of Geography and GeoInformation Science, George Mason University, November 2010.

7. **Mansoor Alicherry**, Thesis title: "*A Distributed Policy Enforcement Architecture for Mobile Ad Hoc Networks*",
   Computer Science Department, Columbia University, October 2010.
8. **Min Xu**, Thesis title: "*Session-aware RBAC Administration, Delegation, and Enforcement with XACML*",
   Computer Science Department, George Mason University, April 2010.

Post-Doctoral Researchers

Konstantinos Kolias (August 2014 - 2019)
Daniel Fleck (August 2012 - 2019)
Nelson Nazzicari (August 2010 - September 2011)
Arnur Tokhtabayev (May 2011 - November 2012)

Current Ph.D. Students

Full Time Ph.D. Students

1. Kene Nwondo (Auguest 2020 - present)

Graduated Ph.D. Students

Ryan E. Johnson (January 2011 – December 2019)
- Thesis title: "*Automatic Program State Exploration Techniques for Security Analysis of Android Apps*"
- Post-graduation: Director of Research, Kryptowire LLC
- Currently: Director of Research, Kryptowire LLC

Rahul Murmuria (January 2011 – September 2017)
- Thesis title: "*Modeling User Behavior on Smartphones*"
- Post-graduation: Data Scientist, AppZen
- Currently: Senior Data Scientist at Sensor Tower

Mohamed Elsabagh (September 2012 – June 2017)
- Thesis title: " *Protection from Within: Runtime Hardening Techniques for COTS Binaries*"
- Post-graduation: Kryptowire LLC
- Currently: Kryptowire LLC

Charles Smutz (January 2009 - August 2016) (part-time Ph.D.)
- Thesis title: "*Countering Malicious Documents and Adversarial Learning*"
- Post-graduation: Sandia National Laboratories
- Currently: Sandia National Laboratories

Sharath Hiremagalore (September 2009 - August 2015)
- Thesis title: "*Zero-Day Attack Detection Using Collaborative and Transduction-Based Anomaly Detectors*"
- Post-graduation: Verisign
- Currently: Verisign

Fengwei Zhang (September 2011 - April 2015)
- Thesis title: "*Using Isolated Execution Environments for Securing Systems*"
- Post-graduation: Assistant Professor at Wayne State University
- Currently: Assistant Professor at Wayne State University

Quan Jia (September 2008 - January 2014)
- Thesis title: "*Mitigating Denial-of-Service Attacks in Contested Network Environments*"
- Post-graduation: MicroStrategy
- Currently: MicroStrategy

Zhaohui Wang (September 2008 - December 2012)
- Thesis title: "*Securing Smart Mobile Devices: A Data-Centric Approach*"
- Post-graduation: N/A
- Currently: N/A

Jiang Wang (January 2008 - July 2011)
- Thesis title: "*Hardware-Assisted Protection and Isolation*"
- Post-graduation: Riverbed Technology Inc.
- Currently: Riverbed Technology Inc.

Graduated MSc. Students
1. Charalampos Andrianakis (September 2008 - September 2011)
2. Rhandi Martin (January 2009 - January 2011)
3. Spyridon Panagiotopoulos (September 2009 - December 2011)
4. Chen Liang (September 2009 - December 2011)

Service at George Mason University

Volgenau School of Engineering, Academic Director, M.S. in Management of Secure Information Systems Program, School of Management (March 2014 - 2018)
Computer Science Department, ISA Admissions & Policy Committee (September 2008 - present)
Computer Science Department, Security Recruiting Committee (September 2010 - July 2011)
Computer Science Department, APR ISA Committee (September 2010 - July 2011)
USENIX Association Campus Representative (2010 - present)
Faculty Advisor, undergraduate student group: **GMU ECHO** (Electrical & Computer Hacking Organization) (September 2009 - 2013)
Faculty Advisor, graduate student group: **GMU Information Security Association** (November 2007 - November 2009)

Teaching Experience

Virginia Tech

(Scores indicate mean course quality rating from student survey out of a maximum of 6.0)

Fall 2021: ECE/CS 5560, Fundamentals of Information Security (**71 students**)

Spring 2020: ECE/CS 5584, Nework Security **(18 students, Instr. :5.33, Class: 5.33)**

Fall 2020: ECE 4614, Telecommunication Networks **(33 students, Instr. :5.47, Class: 5.32)**

George Mason University

(Scores indicate mean course quality rating from student survey out of a maximum of 5.0)

Spring 2019: ISA 564/CS 499, Cyber Security Laboratory

Fall 2018: ISA 673, Operating Systems' Security

Fall 2017: CS 468, Secure Programing and Systems

Fall 2016: ISA 564, CS 499, Cyber Security Laboratory **(14 students, Instr.: 4.43, Class: 4.86)**

Fall 2015: ISA 564/CS 499, Cyber Security Laboratory **(19 students, Instr.: 4.13, Class: 4.31)**

Spring 2015: MSEC510, Foundations of Cyber Security **(31 students, Instr.: 4.39, Class: 4.33)**

Spring 2015: MSEC 650, Seminar: Enterprise Security Case St. **(18 students, Instr.: 4.36, Class: 4.27)**

Spring 2015: MSEC 720, Capstone Project Mgmt. Secure Info **(18 students, Instr.: 4.22, Class: 3.72)**

Spring 2015: ISA 785, Research in Digital Forensics **(13 students, no evaluation)**

Spring 2014: MSEC 511, Enterprise Security Practices **(21 students, Instr.: 4.22, Class: 3.72)**

Spring 2013: MSEC 642, Enterprise Security Technology **(27 students, Instr.: 4.55, Class: 4.54)**

Spring 2013: MSEC 511, Enterprise Security Practices **(21 students, Instr.: 4.67, Class: 4.50)**

Spring 2013: ISA 673, Operating Systems Security **(31 students, Instr.: 4.53, Class: 4.55)**

Fall 2012: ISA 674, Intrusion Detection **(29 students, Instr.: 4.71, Class: 4.71)**

Spring 2012 - Fall 2007. Weighted Average, Instructor: 4.66, Class: 4.53

Spring 2012: ISA 673, Operating Systems' Security **(36 students, Instr.: 4.71, Class: 4.48)**
Spring 2012: MSEC 511, Enterprise Security Practices **(30 students, Instr.: 4.93, Class: 4.77)**

Fall 2011: ISA 785, Research in Digital Forensics **(29 students, Instr.: 4.91, Class: 4.91)**
Fall 2010: ISA 862, Models for Computer Security (**23 students, Instr.: 4.89, Class: 4.84)**

Spring 2010: ISA 673, Operating Systems Security (**28 students, Instr.: 4.46, Class: 4.58)**

Fall 2009: CS 571, Operating Systems **(40 students, Instr.: 4.58, Class: 4.21)**

Spring 2009: ISA 564, Security Laboratory **(46 students, Instr.: 4.45, Class: 4.42)**

Fall 2008: ISA 656, Network Security **(28 students, Instr.: 4.81, Class: 4.69)**

Spring 2008: IT 862, Models for Computer Security **(29 students, Instr.: 4.38, Class: 4.25)**

Spring 2008: ISA 656, Network Security **(32 students, Instr.: 4.64, Class: 4.46)**

Fall 2007: ISA 656, Network Security **(30 students, Instr.: 4.68, Class: 4.50)**

Awarded Support for Research and Teaching (Gifts and Grants)

Total: ~$25,933,500 Total as PI: **~$22,082,500**

1. **PI DARPA STTR Twisted, Phase I, $67,500, 02/01/2021 – 11/30/2021** "Highly assured trusted information storage device" (with Trusted Science and Technology Inc.)

2. **PI DARPA WASH, ~$6,700,000, 03/01/2018 – 02/30/2022 via Kryptowire LLC** "SHINE: Sensing Health Innovatively Using Novel Empiricism"

3. **PI NIST, $473,632, 09/01/2016 - 08/31/2019** "Towards Measuring Security for IoT"

4. **Co-PI NSF, $299,935, 09/01/2016 - 08/31/2019** "City and County Cross Jurisdiction Cybersecurity Collaboration Capacity Building" (with J.P. Auffret)
5. **PI DARPA LADS (sub to PFP Cyber), $1,454,051, 05/01/2016 - 04/30/2020** "Enhanced Cyber Defense by Leveraging Involuntary Analog Emissions" (with J.P. Auffret)
6. **PI DARPA XD3, $4,433,701, 04/12/2016 - 06/30/2020** "Democratizing DDoS Defenses Using

   Secure Indirection Networks" (GMU-lead team with Columbia University, Penn. State, and BAE Systems, GMU portion **$1,529,742** without options) (with Dan Fleck)
7. **PI DAPRA XD3 (sub to Vencore, Inc), $944,150, 04/20/2016 - 04/19/2019** "Lookout - for the DARPA Extreme DDoS Defense- TA3" (with Dan Fleck)

8. **PI, Korea Agency for Defense Development, $267,682, 06/02/2014 - 02/15/2016,** "Technical consulting on the test and evaluation methodology for cyber-security technologies" (with J.P. Auffret)
9. **PI, NSF, $174,900, 09/01/2014 - 08/31/2017,** "TWC: TTP Option: Small: Collaborative: Scalable Techniques for Better Situational Awareness: Algorithmic Frameworks and Large-Scale Empirical Analyses" (with Fabian Monrose, UNC)
10. **Co-PI, DARPA (sub to Invincea Labs), $360,753, 01/15/2014 - 03/30/2015,** "TAPIO: Targeted Attack Premonition using Integrated Operational data sources" (with Dan Fleck)
11. **PI, NSF, $484,857, 08/01/2013 - 07/30/2016,** "Bridging the Cybersecurity Leadership Gap: Assessment, Competencies and Capacity Building" (with J.P. Auffret)
12. **PI, DHS/Purdue, $486,691, 07/01/2013 - 06/30/2016,** "Analysis of Mobile Application Communications Using GUI & Data Instrumentation"
13. **Co-PI, DHS, $256,000, 09/20/2012 - 08/31/2017,** "Graduate Fellowship Training for Homeland Security" (with Duminda Wijesekera and Damon McCoy)
14. **co-PI Google Research Award, $75,000, 06/2013** (with Damon McCoy)
15. **co-PI NSF II-New, $547,000 09/2012-08/2013,** "An Experimental Infrastructure for Cross-Domain Research in Wireless Computing, Cybersecurity and Data" (with Robert Simon, Daniel Barbara and Brian Mark)
16. **PI (GMU), DARPA MRC, $750,363 09/2011 - 01/2016,** "**MEERKATS**: **M**aintaining **E**nterpris**E** **R**esiliency via **K**aleidoscopic **A**daptation & **T**ransformation of **S**oftware **S**ervices",
(Part of team that includes Columbia University and Symantec Corp. total budget: $6,619,270) (with Fei Li)
17. **PI, DARPA Transformative Applications/Aterrasys, $511,323 08/24/2011 - 08/24/2012,** "Securing Android Mobile Devices"
18. **PI, Army Research Office (ARO), DURIP $205,983 06/15/2011 - 06/14/2012,** "A VPN Proxy Cloud for Detecting HTTP & VoIP Malware" (with Anup Ghosh)
19. **PI, IARPA, $2,169,506 08/02/2010 - 05/31/2014,** "Securely Taking on New Executable Software of Uncertain Provenance (STONESOUP) Program" (with Anup Ghosh)
20. **PI, DARPA, $1,527,225 07/01/2010 - 06/30/2014**, "CyNomix: Detecting Zero-Day Malware by Generating Behavioral Cyber Genome Sequences" (with Huzefa Rangwala)
21. **PI, NIST/DARPA, $653,780 (+$300,000 Supplement) 08/01/2010 - 07/31/2013,** "Securing Android Smart-Phones via Automated Testing and Certified Communications" (with Anup Ghosh)
22. **co-PI, NIST, $431,902 07/01/2010 - 06/30/2013,** "Building Policies to Control Virtual Environments using the Policy Machine" (with Duminda Wijesekera)
23. **co-PI, DHS, $368,923/$980,000 08/27/2010 - 05/31/2011 (2010),** "ATHENA-Yukon Project" (with Anup Ghosh)
24. **co-PI, Secure Command, LLC $32,797 09/01/2010 - 03/31/2011,** "Enforcing Hardware-Assisted Integrity & Trust for Commodity Operating Systems" (with Kun Sun)

25. **PI, NSF, $239,884 09/2009-08/2011**, "TC: Small: Collaborative Research: Scalable Malware Analysis Using Lightweight Virtualization" (with Fabian Monrose)
26. **PI, Army Research Office (ARO), $342,400 09/2009-08/2011**, STTR Phase II: "Automatic Identification & Mitigation of Unauthorized Information Leaking from Enterprise Networks" (with Sushil Jajodia)
27. **co-PI, DARPA, $291,000 09/2009-08/2010**, "An Architecture for Providing High Assurance of Untrusted Applications on Wireless Handheld Devices" (with Anup Ghosh).
28. **co-PI, BAE Systems/DARPA, $59,875 1/1/09 - 09/11/2009,** "National Cyber Range" (with Anup Ghosh)
29. **PI, Google Inc: Research gift, $90,000 03/09,** (with Fabian Monrose)
30. **co-PI, AFOSR, $250,675 08/2009-08/2010**, DURIP: "A Laboratory for Large-Scale Testing of Self-Healing" (with Anup Ghosh)
31. **co-PI, Princeton University/DARPA, $84,937 8/16/08 - 8/31/09,** "Parallelizing Legacy Binary Code for Multi-Core Architectures via Extraction of Self-Similarity" (with Michael Locasto)
32. **co-PI, Army Research Office (ARO), DURIP $150,000 07/2009-07/2009**, "A Laboratory for Proactively Preventing Phishing and Malcode Attacks Using Web Crawlers", (with Sushil Jajodia and Anup Ghosh)
33. **co-PI, DHS/I3P Dartmouth College, $60,000 11/2009,** "Securing the Railway IT Infrastructure", (with Michael Locasto and Duminda Wijesekera)
34. **co-PI, AFOSR, $670,499 07/2009-07/2011**, "Secure Composition of Networked Systems Based on User Tasks and Organizational Policy" (with Duminda Wijesekera and Sushil Jajodia).
35. **co-PI, DARPA/BAE Systems, $50,000 1/1/09 - 6/30/09** "National Cyber Range" (with Anup Ghosh)
36. **PI DHS/I3P Dartmouth College: $150,000 8/10/08 - 8/9/09** "Open Taint: Flexible and Automatic Dataflow Tagging and Control for User-Level Programs" (with Michael Locasto)
37. **co-PI, Google Inc: Research gift, $25,000 03/08,** (with Steven M. Bellovin)
38. **co-PI, Secure Command, LLC: $50,000 9/19/08 - 3/18/09** "STTR: Fingerprinting Network Traffic" (with Sushil Jajodia)

Research Experience

Computer Science department, Columbia University,
Fu Foundation School of Engineering & Applied Science, New York, NY.
Research Assistant (Fall 2003 - Summer 2007).
Design and Implementation of protection mechanisms against DDoS Attacks using Overlay networks. |
NSL Web page has more info on SOS/WEBSOS project.

Electrical Engineering department, Columbia University,
Fu Foundation School of Engineering & Applied Science, New York, NY.
Research Assistant (Spring 2002 - Fall 2003).
Design and implementation of a novel peer to peer client/server protocol in Java.
Performed Internet experiments using up to 180 concurrent nodes in various locations around the world.

European Union program TIDE/RISE for home networks application.
Development of robust home network applications for a controlled medical environment.

General Secretariat of Research and Technology of Greece.
Design and implementation of Industrial network for the Kopais industry as a part of a program from the
General Secretariat of Research and Technology of Greece.

Prior to 2001 Work Experience

01/1999 - 12/2000: Network Administrator, University of Athens
03/1997 - 07/1998: Network Administrator. Westnet S.A.
09/1994 - 09/1997: University of Patras, Network Administrator

Academic Honors, Fellowships

Outstanding Research Award: 2016 Department of Computer Science, George Mason University.

IEEE Reliability Society Engineer of the Year Award (2012) - awarded January 2013.

Mason Masters In Secure Information Systems Outstanding Faculty of the Year Award (2013, 2014).

Mason Emerging Researcher/Scholar/Creator award: 2012 George Mason University (one out of three awards for 2012-2013).

Outstanding Research Award: 2010 Department of Computer Science, George Mason University.
Dissertation with Distinction Award: 2007 Computer Science Department, Columbia University.

CS Service Award: 2006 Computer Science Department, Columbia University.

Preceptor: Columbia University Fellow Spring 2004 & Fall 2005.

Best Teaching Assistant Award: Spring 2002, Columbia University.

Scholarship: from the graduate program of Algorithms, Logic & Computation for the first two years of study (1998-2000).

Greek National Fellowship Institution award: for being the second (2/180) for the first and third years of undergraduate study.

Professional References

Professor Angelos D. Keromytis, John H. Weitnauer, Jr. Chair, and Georgia Research Alliance (GRA) Eminent Scholar at the Georgia Institute of Technology
Contact: angelos@gatech.edu
Klaus 3362, 266 Ferst Dr NW, Atlanta, GA 30332
https://www.ece.gatech.edu/faculty-staff-directory/angelos-d-keromytis

Professor Fabian Monrose, Kenan Distinguished Professor in the Computer Science University of North Carolina at Chapel Hill
Contact: fabian@cs.unc.edu
3175 Sitterson Hall, UNC-Chapel Hill, NC 27599-3175
https://www.cs.unc.edu/~fabian/web.html

Professor Salvatore J. Stolfo, Computer Science Department, Columbia University
Contact: sal@columbia.edu
606 CEPSR, Mail Code 0401, 530 West 120th Street, New York, NY 10027
https://salvatorestolfo.com/

Professor Jonathan M. Smith, Olga and Alberico Pompa Professor of Engineering and Applied Science, Professor of Computer and Information Science at the University of Pennsylvania
Contact: jms@cis.upenn.edu
604 Levine Hall, CIS Dept., 3330 Walnut St. Philadelphia, PA 19104-6389
https://www.cis.upenn.edu/~jms/

Professor Steven M. Bellovin, Computer Science Department, Columbia University
Contact: smb@cs.columbia.edu
454 Computer Science Building, 500 West 120th St, M.C. 0401, New York, NY 10027-7003
https://www.cs.columbia.edu/~smb/

APPENDIX C

Robustness Testing of the Microsoft Win32 API

Charles P. Shelton
ECE Department & ICES
Carnegie Mellon University
Pittsburgh, PA, USA
cshelton@cmu.edu

Philip Koopman
ECE Department
Carnegie Mellon University
Pittsburgh, PA, USA
koopman@cmu.edu

Kobey DeVale
ECE Department
Carnegie Mellon University
Pittsburgh, PA, USA
kdevale@ece.cmu.edu

Abstract

Although Microsoft Windows is being deployed in mission-critical applications, little quantitative data has been published about its robustness. We present the results of executing over two million Ballista-generated exception handling tests across 237 functions and system calls involving six Windows variants, as well as similar tests conducted on the Linux operating system. Windows 95, Windows 98, and Windows CE were found to be vulnerable to complete system crashes caused by very simple C programs for several different functions. No system crashes were observed on Windows NT, Windows 2000, and Linux. Linux was significantly more graceful at handling exceptions from system calls in a program-recoverable manner than Windows NT and Windows 2000, but those Windows variants were more robust than Linux (with glibc) at handling C library exceptions. While the choice of operating systems cannot be made solely on the basis of one set of tests, it is hoped that such results will form a starting point for comparing dependability across heterogeneous platforms.

1. Introduction

Different versions of the Microsoft Windows operating system (OS) are becoming popular for mission- and safety-critical applications. The Windows 95/98 OS family is the dominant OS used in personal computer systems, and Windows NT 4.0 has become increasingly popular in business applications. The United States Navy has adopted Windows NT as the official OS to be incorporated into onboard computer systems [15]. Windows CE and Windows NT Embedded are new alternatives for embedded operating systems. Thus, there is considerable market and economic pressure to adopt Windows systems for critical applications.

Unfortunately, Windows operating systems have acquired a general reputation of being less dependable than Unix-based operating systems. In particular, the infamous "Blue Screen Of Death" that is displayed as a result of Windows system crashes is perceived by many as being far more prevalent than the equivalent kernel panics of Unix operating systems. Additionally, it is a common (although meagerly documented) experience that Windows systems need to be rebooted more often than Unix systems. However, there is little if any quantitative data published on the dependability of Windows, and no objective way to predict whether the impending move to Windows 2000 will actually improve dependability over either Windows 98 or Windows NT.

Beyond the dependability of Windows itself, the comparative dependability of Windows and Unix-based systems such as Linux has become a recurring theme of discussion in the media and Internet forums. While the most that can usually be quantified is mean time between reboots (anecdotally, Unix systems are generally said to operate longer between reboots than Windows NT), issues such as system administration, machine usage, the behavior of application programs, and even the stability of underlying hardware typically make such comparisons problematic. It would be useful to have a comparison of reliability between Windows and Unix systems based on direct, reproducible measurements on a reasonably level playing field.

The success of many critical systems requires dependable operation, and a significant component of system dependability can be a robust operating system. (Robustness is formally defined as the degree to which a software component functions correctly in the presence of exceptional inputs or stressful environmental conditions [6].) Of particular concern is the behavior of the system when confronted with exceptional operating conditions and consequent ex-

ceptional data values. Because these instances are by definition not within the scope of designed operation, it is crucial that the system as a whole, and the OS in particular, react gracefully to prevent compromising critical operating requirements. (Some systems, such as clustered web servers, can be architected to withstand single-node failures; however there are many critical systems in the embedded computing world and mission-critical systems on desktops which cannot afford such redundancy, and which require highly dependable individual nodes.)

This paper presents a quantitative comparison of the vulnerability of six different versions of the Windows Win32 Application Programming Interface (API) to robustness failures caused by exceptional function or system call parameter values. These results are compared to the results of similarly testing Linux for exception handling robustness.

Exception handling tests are performed using the Ballista robustness testing harness [1] for both Windows and Linux. In order to perform a reasonable Windows-to-Linux comparison, 237 calls were selected for testing from the Win32 API, and matched with 183 calls of comparable functionality from the Linux API. Of these calls, 94 were C library functions that were tested with identical test cases in both APIs, with the balance of calls being system calls. Beyond C library functions, the calls selected for testing were common services used by many application programs such as memory management, file and directory system management, input/output (I/O), and process execution/control. The results are reported in groups rather than as individual functions to provide a reasonable basis for comparison in those areas where the APIs differ in the number and type of calls provided.

2. Background

The Ballista software testing methodology has been described in detail elsewhere [3], [9] and is publicly available as an Internet-based testing service [1] involving a central testing server and a portable testing client that was ported to Windows NT and Windows CE for this research. Thus, only a brief summary of Ballista testing operation will be given.

The Ballista testing methodology is a combination of software testing and fault injection approaches. Specifically selected exceptional values (selected via typical software testing strategies) are used to inject faults into a system via an API. For testing an OS, this involves selecting a set of functions and system calls to test, with each such Module under Test (MuT) being exercised in turn until a desired portion of the API is tested. Parameter test values are distinct values for a parameter of a certain data type that are randomly drawn from pools of predefined tests, with a separate pool defined for each data type being tested. These pools of values contain exceptional as well as non-exceptional cases to avoid successful exception handling on one parameter from masking the potential effects of unsuccessful exception handling on some other parameter value. Each test case (the execution of a single MuT with a single test value selected for each required parameter in the call) is executed as a separate task to minimize the occurrence of cross-test interference. A single Ballista test case involves selecting a set of test values, executing constructors associated with those test values to initialize essential system state, executing a call to the MuT with the selected test values in its parameter list, measuring whether the MuT behaves in a robust manner in that situation, and cleaning up any lingering system state in preparation for the next test (including freeing memory and deleting temporary files).

Ballista testing looks only for non-robust responses from software, and does not test for correct functionality. This, combined with a data type-based testing strategy, rather than a functional testing strategy, results in a highly scalable testing approach in which the effort spent on test development tends to grow sub-linearly with the number of MuTs to be tested. An additional property of Ballista testing results is that in practice they have proven to be highly repeatable. Virtually all test results reproduce the same robustness problems every time a brief single-test program representing a single test case is executed.

Ballista uses the CRASH scale [9] to measure robust or non-robust responses from MuTs. CRASH is an acronym for the different robustness failures that can occur. In **C**atastrophic failures, the most severe robustness failure type, the application causes a complete system crash that requires an OS reboot for recovery. In **R**estart failures, the application enters a state where it "hangs" and will not continue normal operation, requiring an application restart for recovery. **A**bort failures are an abnormal termination of an application task as the result of a signal or thrown exception that is not specific enough to constitute a recoverable error condition unless the task elects (or is forced by default) to terminate and restart. **S**ilent failures occur when a function or call is performed with invalid parameter values, but the system reports that it was completed successfully instead of returning an error indication. Finally, **H**indering failures report an incorrect error indication such as the wrong error reporting code. Ballista can automatically detect Catastrophic, Restart, and Abort failures; Silent failures and Hindering failures currently can be detected in only some situations, and require manual analysis.

Earlier Ballista publications (e.g., [3], [8], [9]) describe the software testing and fault injection heritage of this approach. Ballista can be thought of as using software testing principles to perform fault injection at the API level instead

of the source code or object code level. The most closely related current research effort is the work done at Reliable Software Technologies on testing Windows NT [4], [5] in light of Ballista results on Unix systems. That work focuses on a broad coverage of functions for a single OS version with relatively simple testing values. Nonetheless, their results found many Abort-type failures in Windows NT, and a few Catastrophic failures that were caused by very complex execution sequences that could not be isolated for bug-reporting purposes. Other recent related work is the Fuzz project at the University of Wisconsin [12], [13], which has concentrated on Unix systems. There does not appear to be any previously published work that performs testing-oriented dependability comparisons of multiple Windows versions, nor comparisons of Windows to Unix robustness.

3. Implementation

The existing Ballista testing system ran only on Unix systems. Thus, testing Windows required porting the client-side testing harness to Windows as well as creating Windows-specific test values and an inter-API comparison methodology.

3. 1. Porting to Windows Desktop Systems

Porting the Ballista testing client software to the Windows platform faced many difficulties, chief among them the fact that Windows has no simple analog to the fork() system call implemented on POSIX systems (POSIX [7] is the standard for Unix). Thus it is more difficult to spawn a child process for each test case being executed. To overcome this, the Windows version of the Ballista test harness creates a memory-mapped file for each test case, writes data for that particular test case's parameters to this file, and then spawns the testing process. The testing process retrieves the data for the current test case from the memory location created by the calling process, and reports results for the test to that same memory location.

The Win32 API uses a thrown-exception error reporting model in addition to the error return code model (using the POSIX "errno" variable or the Win32 GetLastError() function) used by the POSIX API. While on POSIX systems Abort failures can be detected by simply monitoring the system for the occurrence of signals (most often `SIGSEGV` or `SIGBUS`), in Windows systems there are both legitimate and non-robust occurrences of thrown error reporting conditions. The Win32 API documentation [11], [14] does not provide sufficient information to make a per-function list of permissible and non-permissible thrown exceptions. For Windows testing, the Ballista test harness intercepted all integer and string exception values, and to be more than fair in evaluation, assumed that all such exceptions were valid and recoverable. In normal operation, any unrecoverable exceptions trigger the Windows top-level exception filter and display an "Application Error" message window before terminating the program. We disabled this exception filter and replaced it with code that would record such an unrecoverable exception as an Abort failure. (This technique could in fact be used to improve the robustness of an application program, but only by restarting abnormally terminated tasks. That approach might be sufficiently robust for many users, but is considered to be non-robust at the application level by most of the critical-system designers we have had discussions with.)

There were additional challenges involved in porting the Ballista testing client to a Windows environment, such as obtaining a remote procedure call (RPC) package that was compatible with the Unix-based Ballista testing server's RPC implementation. Most UNIX systems use ONC RPC, but Windows only supports DCE RPC, so a third party ONC RPC Windows client had to be used. Most porting issues were related to differing OS interface architectures, and were not fundamental to the Ballista approach.

Because many Win32 calls have four or more parameters, a very large number of test cases could be generated without exhausting all potential combinations of test values for a single MuT. Therefore, testing was capped at 5000 randomly selected test cases per MuT. 72 Windows MuTs and 34 POSIX MuTs were capped at 5000 tests each (per OS) in this manner. All other MuTs performed exhaustive testing of all combinations with fewer than 5000 tests. In order to fairly compare the desktop Windows variants, the same pseudorandom sampling of test cases was performed in the same order for each system call or C function tested across the different Windows variants. Previous findings have indicated that this random sampling gives accurate results when compared to exhaustive testing of all combinations [9].

The Win32 and POSIX APIs use different data types. However, most of the Windows data types required were minor specializations of fairly generic C data types. In those cases, the same test values used in POSIX were simply used for testing Windows. The only major data type for which new test values had to be created for testing Windows was the HANDLE type. The tests for this type were largely created by inheriting tests from existing types and adding test cases in the same general vein as existing data type tests. Overall, the data values used for testing were selected based on experience with previous Ballista testing and a general background knowledge from the software testing literature [2].

3. 2. Porting to Windows CE

The Ballista client for Windows NT does not work on the Windows CE platform because Windows CE is designed to be an embedded operating system that runs on specialized hardware for consumer electronics and mission-critical systems. These systems have tighter memory constraints than a normal desktop PC. Also, Windows CE programs must be compiled and linked for specific hardware using tools that run on Windows NT, and then downloaded to the Windows CE device.

To overcome this problem, the Ballista client was split into two components: the test generation and reporting functions that run on a Windows NT PC, and the test execution and control functions that run on the target Windows CE platform. For each system call or function tested, the test execution and control portion is compiled on the PC and downloaded to the Windows CE machine via a serial port connection. The test generation component running on the PC initiates each test case by starting the test execution process on the target and passing the parameter list via the command line arguments.

Windows CE provides a remote API that allows Windows NT applications to communicate with the Windows CE target using file I/O and process creation, but does not provide mechanisms for process synchronization or control. Therefore, the test execution component running on the target must create another process that actually runs the test and records the result in the target's file system. The NT process must remain idle and wait for this file to appear on the target to get the results of the current test case and report them. Unfortunately this means tests are several orders of magnitude slower than tests run on the other Windows OS versions, taking five to ten seconds per test case.

Error classification was also a problem on Windows CE. Windows CE does not support the normal C++ `try/catch` exception handling scheme, so we had to use the Win32 structured exception handling constructs, `__try/__except` and `__try/__finally`. We did not use these on the other Windows platforms because the Microsoft documentation [11] recommends using C++ `try/catch` whenever possible, and states that the two exception handling methods are mutually exclusive.

The exceptions that we observed on Windows CE appeared to be analogous to the signals thrown in POSIX systems. For example, the exception `EXCEPTION_ACCESS_VIOLATION` thrown in Windows CE is comparable to a `SIGSEGV` signal thrown in UNIX. Therefore, we classified these exceptions as abort failures. The only exceptions observed were `EXCEPTION_ACCESS_VIOLATION`, `EXCEPTION_DATATYPE_MISALIGNMENT`, and `EXCEPTION_STACK_OVERFLOW`.

3. 3. Comparison methodology

Perhaps the greatest challenge in testing Windows systems and then comparing results to Linux was creating a reasonable comparison methodology. While C library functions are identical on both systems, the system calls have different functionality, different numbers of parameters, and somewhat different data types. However, the Ballista techniques of basing tests on data types and of normalized failure rate reporting were used to create an arguably fair comparison of exception handling test results.

Basing tests on data types rather than MuT functionality permits comparing APIs with similar functionality but different interfaces. Because the data type test definitions are nearly identical for both Windows and Linux, the same general tests in the same general proportions are being run regardless of functionality. Of course there is always the possibility of accidental bias. But, because the tests were originally developed specifically to find problems with POSIX systems by students who had no Windows programming experience, if anything the tests would be biased toward finding problems on the previous testing target of POSIX rather than specifically stressing Windows features.

Normalized failure rate data was used to permit comparison among different interfaces to similar functionality between Windows and Linux. Normalization is performed by computing the robustness failure rate on a per-MuT basis (number of test cases failed divided by number of test cases executed for each individual MuT). Then, the MuTs are grouped into comparable classes by functionality, such as all MuTs that perform memory management. The individual failure rates within each such group are averaged with uniform weights to provide a group failure rate, permitting relative comparisons among groups for all OS implementations. As an example, the I/O Primitives group consists of `{close dup dup2 fcntl fdatasync fsync lseek pipe read write}` for POSIX and `{AttachThreadInput CloseHandle DuplicateHandle FlushFileBuffers GetStdHandle LockFile LockFileEx ReadFile ReadFileEx SetFilePointer SetStdHandle UnlockFile UnlockFileEx WriteFile WriteFileEx}` for Win32. Robustness failure rates for the I/O Primitives group are computed by averaging the 10 individual failure rates for the POSIX calls, and comparing against the averaged result for the 15 individual Win32 call failure rates for some particular Windows implementation. While even this level of comparison obviously is not perfect, it has the virtue of encompassing the same set of higher-level functionality across two different APIs. For the purposes of achieving generic-level functionality comparisons, calls that did not have an obvious grouping counterpart for both POSIX and Windows were discarded.

Table 1. Robustness Failure rates by Module under Test (MuT) for Windows versions and Linux.

	System Calls Tested	System Calls with Catastrophic Failures	System Calls with Calculated Failure Rates	System Percent Restart Failures by Call	System Percent Abort Failures by Call	C Library Functions Tested	C Library Functions with Catastrophic Failures	C Library Functions with Calculated Failure Rates	C Library Percent Restart Failures by Function	C Library Percent Abort Failures by Function
Linux	91	0	91	0.2%	7.1%	94	0	94	0.8%	34.9%
Windows 95	133	7	126	0.1%	11.6%	94	1	93	0.02%	24.7%
Windows 98	143	5	138	0.1%	13.3%	94	2	92	0.0%	24.6%
Windows 98 SE	143	6	137	0.1%	12.9%	94	1	93	0.0%	25.0%
Windows NT	143	0	143	0.3%	23.5%	94	0	94	0.01%	24.6%
Windows 2000	143	0	143	0.4%	22.7%	94	0	94	0.05%	24.1%
Windows CE	71	10	61	0.1%	13.3%	82 (108)	18 (27)	64	0.0%	14.0%

	Total MuTs (Functions + Calls) Tested	Overall MuTs with Catastrophic Failures	Overall MuTs with Calculated Failure Rates	Overall Percent Restart Failures by MuT	Overall Percent Abort Failures by MuT
Linux	183	0	183	0.5%	21.9%
Windows 95	227	8	219	0.08%	17.2%
Windows 98	237	7	230	0.06%	17.8%
Windows 98 SE	237	7	230	0.06%	17.8%
Windows NT	237	0	237	0.20%	23.9%
Windows 2000	237	0	237	0.23%	23.3%
Windows CE	153 (179)	28 (37)	125	0.04%	13.7%

In all, 3,430 distinct test values incorporated into 37 data types were available for testing POSIX, and 1,073 distinct test values incorporated into 43 data types were available for testing Windows. Given the cap of 5000 tests per MuT, a total of over 148,000 tests were run on each implementation of the C library, plus an additional 380,000 tests on each implementation of the Win32 API compared to 210,000 tests on the Linux system calls.

We did not test any functions in the Graphical Device Interface (GDI) or any Windows device driver specific code. Similarly, although we did not detect any obvious resource "leakage" during testing, we did not specifically target that type of failure mode for testing, nor did we test the systems under heavy loading conditions. While these are clearly potential sources of robustness problems, we elected to limit testing to comparable situations between Windows and Linux, and to restrict results to include only highly repeatable situations to lend confidence to the accuracy of the conclusions.

4. Experimental Results

Ballista robustness testing was performed on the following operating systems on comparable Pentium-class computers with at least 64 megabytes of RAM:

- Windows 95 revision B
- Windows 98 with Service Pack 1 installed
- Windows 98 Second Edition (SE)/Service Pack 1
- Windows NT 4.0 Workstation/Service Pack 5
- Windows 2000 Professional Beta 3 (Build 2031)
- Windows CE 2.11 running on a Hewlett Packard Jornada 820 Handheld PC
- RedHat Linux 6.0 (Kernel version 2.2.5)

The Microsoft Visual C++ compiler (version 6.0) was used for all Windows systems, and the GNU C compiler (version 2.91.66) was used for the Linux system. (Technically the results for C library testing are the result of the GNU development team and not Linux developers, but they are so prevalently used as a pair to implement POSIX functionality with the C binding that this seems a reasonable approach.)

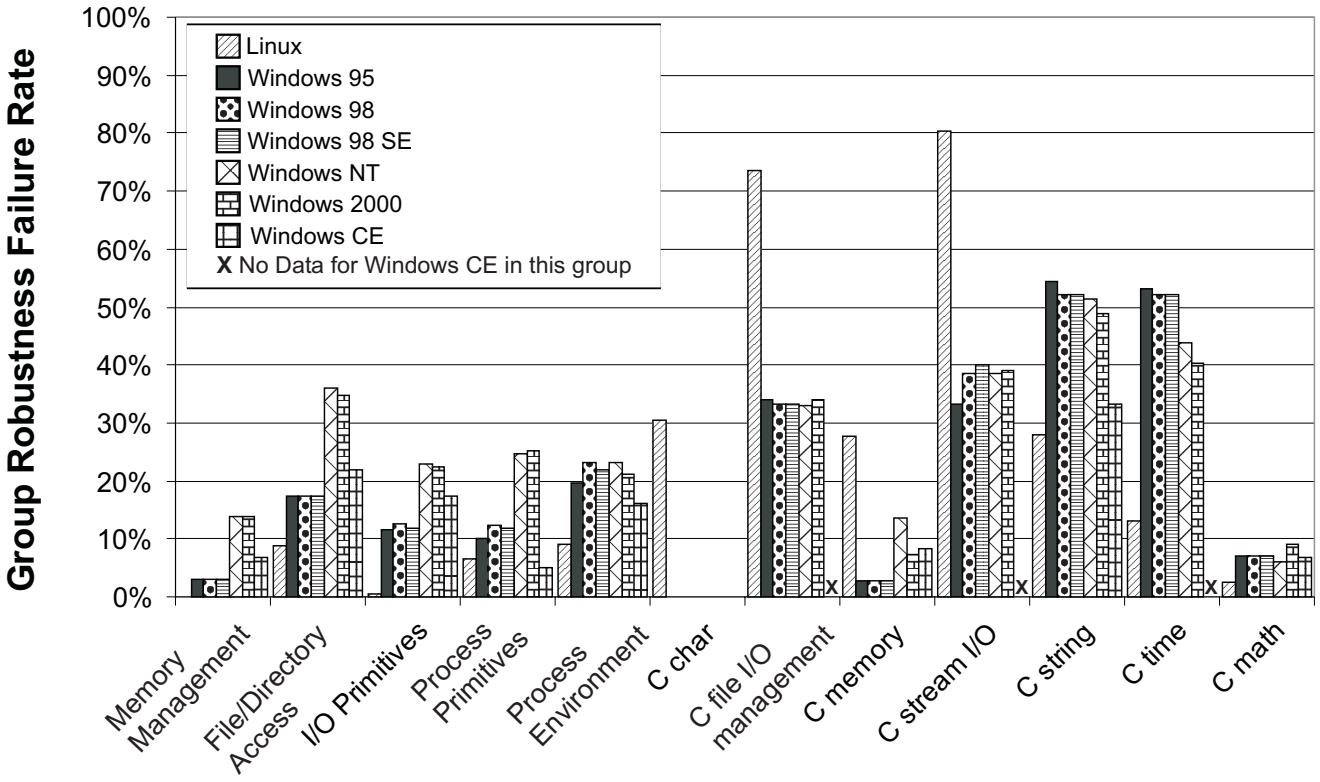

Figure 1. Comparative Windows and Linux robustness failure rates by functional category.

Table 2. Overall robustness failure rates by functional category. Catastrophic failure rates are excluded from numbers, but their presence is indicated by a "*".

	System Calls					C Library						
	Memory Management	File/Directory Access	I/O Primitives	Process Primitives	Process Environment	C char	C file I/O management	C memory	C I/O stream	C string	C time	C math
Linux	0.08%	8.8%	0.4%	6.6%	9.1%	30.4%	73.5%	27.8%	80.4%	28.0%	13.2%	2.6%
Windows 95	*3.0%	*17.3%	*11.6%	*10.2%	*19.6%	0.0%	33.9%	2.7%	*33.2%	54.4%	53.2%	7.1%
Windows 98	3.1%	*17.3%	*12.5%	*12.4%	*23.1%	0.0%	33.3%	2.7%	*38.6%	*52.1%	52.0%	7.1%
Windows 98 SE	3.1%	*17.3%	*11.9%	*11.9%	*22.0%	0.0%	33.3%	2.7%	*39.9%	*52.2%	52.0%	14.8%
Windows NT	13.9%	36.1%	23.0%	24.8%	17.4%	0.0%	32.9%	13.5%	38.7%	51.3%	43.8%	6.2%
Windows 2000	13.9%	34.7%	22.4%	25.3%	13.8%	0.0%	34.0%	7.4%	39.0%	48.9%	40.3%	9.0%
Windows CE	*6.7%	22.0%	17.5%	*5.0%	*16.1%	0.0%	*	8.2%	*	*33.4%	N/A	6.9%

In all, 91 POSIX system calls, 143 Win32 system calls, and 94 C library functions were tested on all desktop operating systems. (10 Win32 system calls were not supported by Windows 95, but were tested on the other desktop Windows platforms.) Because it implements a subset of the Win32 API, only 71 Win32 system calls and 82 C library functions were tested on Windows CE. Table 1 shows the results of robustness testing. The percentages of failures are uniformly weighted averages across all functions tested for that OS. Functions with Catastrophic failures are excluded because the system crash interrupts the testing process, and the set of test cases run for that function is incomplete.

Windows CE gives preferred support to the UNICODE 16-bit character set as opposed to the ASCII 8-bit character set that is used on both UNIX and other Windows platforms. There were 26 C functions that had both an ASCII and a UNICODE implementation. The failure rates for both versions were comparable with the exception of strncpy, which had a Catastrophic failure in the UNICODE version but not in the ASCII version. Since Windows CE uses the UNICODE character set as a default, we only report the failure rates for the UNICODE versions of these C functions. The numbers in parentheses in the Windows CE rows in Table 1 represent the number of functions tested when counting both ASCII and UNICODE functions separately.

In order to compare Windows results to Linux results, the different calls and functions were divided into twelve groupings as shown in Table 2 and Figure 1. These groupings not only serve to permit comparing failure rates across different APIs, but also give a summary of failures for different types of functions. Each failure rate is a uniformly weighted average across all functions tested for that particular OS; the total failure rates give each group's failure rate an even weighting to compensate for the effects caused by different APIs having different numbers of functions to implement each function category. Again, functions with Catastrophic failures are excluded from this calculation. Functions tested on Windows CE in the C file I/O management and the C stream I/O groups had too many functions with Catastrophic failures to report accurate group failure rates; 6 out of 10 in the former and 11 out of 14 in the latter. Windows CE does not support functions in the C time group, so no results for that group are reported.

Table 2 and Figure 1 show that there are significant differences in the robustness failure rates of Linux and Windows, as well as between the Windows 95/98 family and the Windows NT/2000 family of operating systems. Windows CE was unlike either family of desktop Windows variants. (It should be noted that the dominant source of robustness failures is Abort failures, so these results should be interpreted in light of the degree to which those failures affect any particular application.)

Listing 1. A line of code that produces Catastrophic failures on Windows 95, Windows 98 and Windows CE

```
GetThreadContext(GetCurrentThread(),
                                 NULL);
```

Windows 95, Windows 98, and Windows 98 SE exhibited similar failure rates, including a number of functions that caused repeatable Catastrophic system crash failures. Five of the Win32 API system calls: DuplicateHandle(), GetFileInformationByHandle(), GetThreadContext(), MsgWaitForMultipleObjects(), and MsgWaitForMultipleObjectsEx(), plus two C library functions, fwrite() and strncpy(), caused Catastrophic failures for certain test cases in Windows 98. Listing 1 shows a representative test case that has crashed Windows 98 every time it has been run on two different desktop machines, a Windows 95 machine, a Windows 98 laptop computer, and our Windows CE device.

Windows 98 SE had Catastrophic failures in the same five Win32 API system calls as Windows 98, plus another in the CreateThread() call, but eliminated the Catastrophic failure in the C library function fwrite(). Windows 95 had all the Catastrophic failures of Windows 98 except for MsgWaitForMultipleObjectsEx(), which was not implemented in Windows 95. Windows 95 also did not exhibit Catastrophic failures in the C library function strncpy(). Windows 95 did, however, have three additional calls with Catastrophic failures: FileTimeToSystemTime(), HeapCreate(), and ReadProcessMemory().

Windows CE had abort failure rates that did not correspond to either the Windows 95/98 family, or the Windows NT/2000 family, and had significantly more functions with Catastrophic failures than any other OS tested, especially in the C library functions. Windows CE had Catastrophic failures in ten Win32 system calls: CreateThread(), GetThreadContext(), InterlockedDecrement(), InterlockedExchange(), InterlockedIncrement(), MsgWaitForMultipleObjects(), MsgWaitForMultipleObjectsEx(), ReadProcessMemory(), SetThreadContext(), and VirtualAlloc(). Windows CE also had 18 C library functions with Catastrophic failures (27 counting ASCII and UNICODE functions separately), 17 of which failed due to the same invalid C file pointer as a parameter.

For several of the functions with Catastrophic failures we could not isolate the system crash to a single test case. We could repeatedly crash the system by running the entire test harness for these functions, but could not reproduce it when running the test cases independently. These system crashes were probably due to inter-test interference, which indicates that system state was not properly cleaned be-

Table 3. Functions that exhibited Catastrophic failures by OS and function group. A "*" indicates that the failure could not be reproduced outside of the test harness.

	Windows 95	Windows 98	Windows 98 SE	Windows CE
Memory Management				
HeapCreate	X			
VirtualAlloc				X
File/Directory Access				
FileTimeToSystemTime	X			
GetFileInformationByHandle	X	X	X	
Process Primitives				
MsgWaitForMultipleObjects	X	X	X	X
*MsgWaitForMultipleObjectsEx		X	X	X
*ReadProcessMemory	X			X
*CreateThread			X	X
Process Environment				
GetThreadContext	X	X	X	X
*InterlockedDecrement, *InterlockedExchange				X
*InterlockedIncrement, SetThreadContext				X
I/O Primitives				
*DuplicateHandle	X	X	X	
C file I/O management				
clearerr, fclose, fflush				X
_wfreopen, fseek, ftell				X
C I/O stream				
*fwrite	X	X		X
*fread				X
(UNICODE and ASCII) fgetc, *fgets, fprintf,				X
(UNICODE and ASCII) fputc, fputs, fscanf,				X
(UNICODE and ASCII) getc, putc, ungetc				X
C string				
*strncpy		X	X	
(UNICODE) *_tcsncpy				X

tween test cases, even though each test is run in a separate process to minimize this effect. All system calls and functions with Catastrophic failures across all OS's are listed in Table 3 by function group.

Windows NT, Windows 2000, and Linux exhibited no Catastrophic failures during this testing. This is certainly not to say that they cannot be made to crash, but rather that they have reached a different plateau of overall robustness - it is, at a minimum, difficult to find a simple C program that crashes them when run as a single task in user mode. Thus, one can consider that there is some merit to Microsoft's claim that Windows NT is more reliable than Windows 98 (as, for example, stated on their Web site [10]).

Restart failures were relatively rare for all the OS implementations tested. However, they might be a critical problem for any system that assumes fail-fast semantics, including clustered servers that otherwise do not require ultra-high dependability hardware nodes. In, general Restart failures were too infrequent for comparisons to be meaningful.

The classification of Aborts as failures is controversial. In systems in which task termination is acceptable (systems requiring fail-fast operation that can withstand the latency of task restarts), or desirable (debugging scenarios), they may not be considered a problem. However, in some critical and embedded systems that either do not have time to accomplish a task restart or cannot withstand the loss of state information accompanying a task restart, Abort failures can be a significant problem. Our experience in talking with companies that require high levels of field reliability is that Aborts are indeed considered failures for those applications. For other applications that do not share the same philosophy, Abort numbers may not have significant meaning

Given that Abort failures are relevant for some applications, Figure 1 shows that there are striking differences in Abort failure rates across operating systems and functional groupings. For example, Linux has more than a 30% Abort failure rate for C character operations, whereas all the Windows systems have zero percent failure rates (this difference is presumably because Windows does boundary checking on character table-lookup operations). Linux also has higher failure rates on C file I/O management, C stream I/O, and C memory operations. For other groupings Linux has a much lower Abort failure rate. It is interesting to note that the similar code bases for the Windows 95/98 pairing and the Windows NT/2000 pairing show up in relatively similar Abort failure rates. Windows CE generally has lower abort failure rates than Windows NT and Windows 2000, but the significant number of functions that can cause complete system crashes indicates that despite this, Windows CE is less stable than Windows NT/2000.

The issue of Silent failures is a potentially thorny one. Silent failures cannot be measured directly by Ballista be-

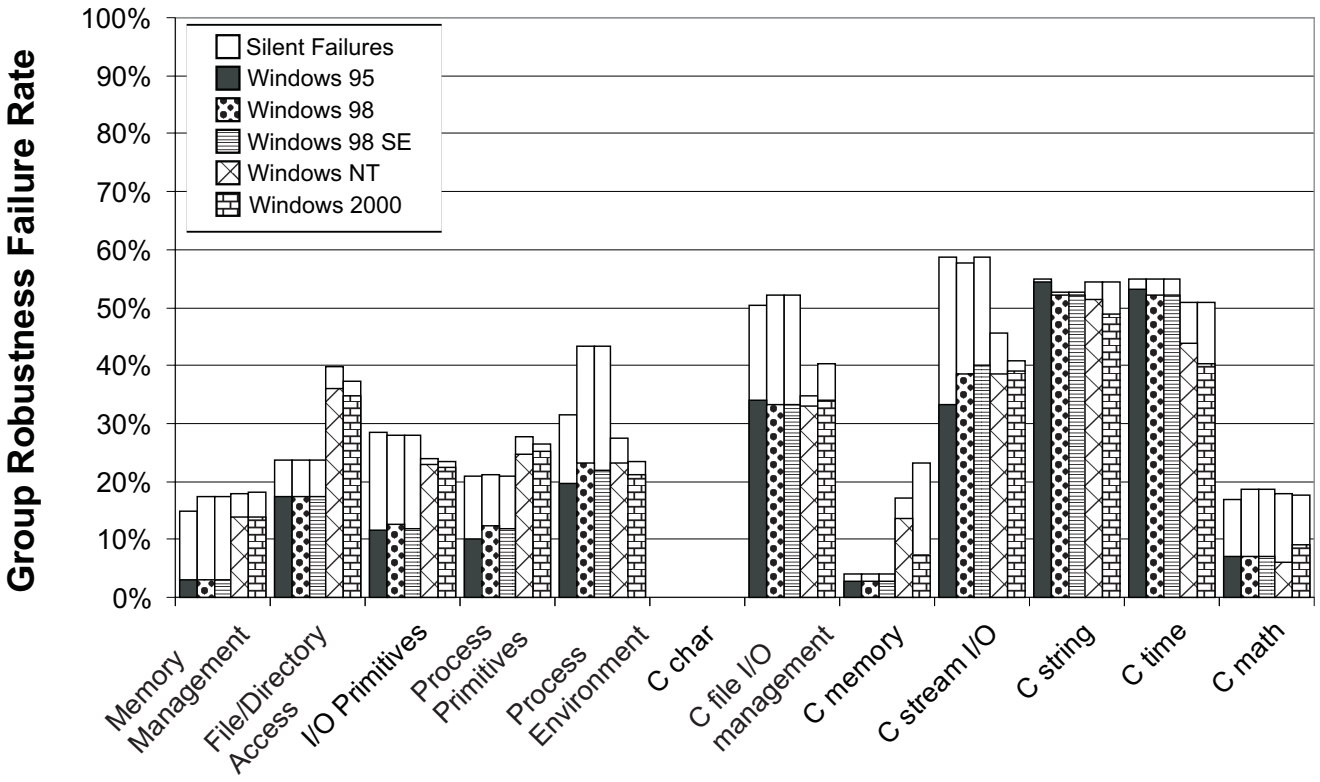

Figure 2. Abort, Restart, and estimated Silent failure rates for Windows desktop operating systems.

cause they involve situations in which there is no observable indication of a failure. (Note: this is not to say they are non-observable in the usual sense of non-activated injected faults that do not affect results. The problem is that there is an exceptional condition that ought to generate observable results to attain robust operation, but does not. As an example, a Silent failure might be a call that reads data from a non-existent file, but returns seemingly valid data bytes with no error indication.)

It is impractical to annotate millions of tests to identify Silent failures. However, we can estimate silent failure rates by voting results across different versions of the same API. Based on previous experience with POSIX [8], we would expect there to be approximately a 10% pass-with-non-exceptional test rate (but, this is a very gross approximation), with the rest of the test cases with a pass with no error reported being Silent failures. If one presumes that the Win32 API is supposed to be identical in exception handling as well as functionality across implementations, if one system reports a pass with no error reported for one particular test case and another system reports a pass with an error or a failure for that identical test case, then we can declare the system that reported no error as having a Silent failure. We wrote a script to automatically vote across identical test cases for each system to generate estimated Silent failure rates. (Note: this analysis does not apply to Linux because it is not an identical API.) Windows CE is not included in this analysis because although the API is similar, it is not identical. Some parameters are not used in Windows CE, and over half of the functions tested on the other Win32 platforms were not supported. Therefore, silent failure rates cannot be reported accurately for Windows CE.

Based on the estimated Silent failures, it seems that the Win32 calls for Windows 95/98/98 SE have a significantly higher Silent failure rate than Windows NT/2000. C library functions vary, with Windows 95/98/98 SE having both higher and lower Silent failure rates than Windows NT/2000 depending on the functional category. Figure 2 shows the overall robustness failure rates for the different Windows variants tested, including these estimated Silent failure rates. Based on these results, it appears that Windows NT and Windows 2000 suffer fewer robustness failures overall than Windows 95/98/98 SE. The only significant exceptions are for the File and Directory Access category as well as the C memory management category, which both suffer from higher Abort failure rates on Windows NT and Windows 2000. (A possible limitation of this approach is that it cannot find instances in which all versions of Windows suffer a Silent failure. This hidden Silent

failure rate may be significant, but quantification is not practical.)

5. Conclusions and Future Work

This work demonstrates that it is possible to compare the robustness of different OS APIs on a relatively level playing field. The use of data type-based testing techniques and normalization of test results by functional groupings enables a detailed comparison of APIs having generally similar capabilities but different interfaces.

Applying the Ballista testing methodology to several Microsoft Windows operating systems revealed a variety of robustness failures. The Windows CE OS and the Windows 95/98/98 SE family of operating systems were clearly vulnerable to robustness failures induced by exceptional parameter values, and could be crashed via a variety of functions. Additionally, the Windows 95/98/98 SE systems had a significant level of Silent failure rates in which exceptional operating situations produced neither abnormal termination nor any other indication of an exception when such an indication was demonstrated possible by other Windows variants. Windows NT and Windows 2000 proved as resistant to system crashes as Linux under these testing conditions, and in most cases had fewer Silent failures than the Windows 95/98/98 SE family (although only a relative comparison was possible; the absolute level of Silent failures is more difficult to determine).

An examination of Abort failures (exceptions too non-specific to be recoverable) and Restart failures (task "hangs") showed differences among the Windows variants and between Windows and Linux. Linux had a significantly lower Abort failure rate in eight out of twelve functional groupings, but was significantly higher in the remaining four. The four groupings for which Linux Abort failures are higher are entirely within the C library, for which the POSIX and Win32 APIs are identical.

Windows CE has abort failure rates comparable to Windows NT and Windows 2000, but has several functions that cause complete system crashes. This makes Windows CE a less attractive alternative for embedded systems, where dependability and reliability are of much higher importance than in desktop PC applications. While abort failures may be recoverable by task restarts, a complete OS crash will more than likely cause complete system failure. It should be noted that many of the catastrophic failures found in Windows CE were traceable to incorrect handling of a single bad parameter value, namely an invalid C file pointer (the actual parameter was a string buffer typecast to a file pointer). It could be argued that since we can trace problem to one underlying cause that we should not penalize Windows CE for seventeen functions that happen to take the same parameter. However, developers who wish to use Windows CE in their systems would have to generate software wrappers for each of the seventeen functions they use to protect against a system crash because they only have access to the interface, not the underlying implementation.

It is also interesting to note that several of the Win32 system calls that crashed on Windows CE also crashed on Windows 95/98/98 SE (some with the exact same parameter values, as in Listing 1), despite the fact that they were developed by different teams within Microsoft and have different code bases. One can speculate that this indicates the underlying causes of these errors may be in the specification rather than the implementation; however the problem may simply be that different programmers tend to make the same sorts of mistakes in similar situations.

While it is not appropriate to make sweeping claims about the dependability of Windows or Linux from these test results alone, a few observations seem warranted by the data presented. The marked difference in finding catastrophic failures in Windows CE and the Windows 95/98/98 SE family compared to the other OS families lends credibility to Microsoft's statement that the Windows NT/2000 systems are more reliable overall. A relative assessment of Linux vs. Windows NT reliability is less clear-cut. Linux seems more robust on system calls, but more susceptible to Abort failures on C library calls (which are actually part of the GNU C compiler suite for Linux) compared to Windows NT.

Future work on Windows testing will include looking for dependability problems caused by heavy load conditions, as well as state- and sequence-dependent failures. In particular, we will attempt to find ways to reproduce the elusive crashes that we have observed to occur in both Windows and Linux outside of the current robustness testing framework.

6. Acknowledgements

This work was supported by Emerson Electric, Lucent Technologies, Asea Brown Boveri (ABB), and Darpa (contract DABT63-96-C-0064). Equipment support was provided by Intel, Compaq and Microsoft. Thanks to Jiantao Pan and Meredith Beveridge for their assistance and to Dan Siewiorek for his continuing guidance.

7. References

[1] Ballista Robustness Testing Service, *http://www.ices.cmu.edu/ballista/index.html*, November 1999.

[2] Beizer, Boris, *Black Box Testing: Techniques for Functional Testing of Software Systems*. John Wiley & Sons, Inc., New York, 1995.

[3] DeVale, J., Koopman, P., Guttendorf, D., "The Ballista Software Robustness Testing Service," *Testing Computer Software Conference*, 1999.

[4] Ghosh, A., Schmid, M., Hill, F., "Wrapping Windows NT Software For Robustness," *29th Fault Tolerant Computing Symposium*, June15-18, 1999.

[5] Ghosh, A., Schmid, M., "An Approach to Testing COTS Software for Robustness to Operating System Exceptions and Errors," *10th International Symposium on Software Reliability Engineering*, November 1-4, 1999.

[6] *IEEE Standard Glossary of Software Engineering Terminology (IEEE Std610.12-1990)*, IEEE Computer Soc., Dec. 10, 1990.

[7] *IEEE Standard for Information Technology - Portable Operating System Interface (POSIX) Part 1: System Application Program Interface (API) Amendment 1: Realtime Extension [C Language]*, IEEE Std 1003.1b-1993, 1994.

[8] Koopman, P., DeVale, J., "Comparing the Robustness of POSIX Operating Systems," *29th Fault Tolerant Computing Symposium*, June 15-18, 1999.

[9] Kropp, N., Koopman, P. & Siewiorek, D., "Automated Robustness Testing of Off-the_Shelf Software Components," *28th Fault Tolerant Computing Symposium*, June 23-25, 1998.

[10] Microsoft Corp., *Choosing the Best Windows Desktop Platform For Large and Medium-Sized Businesses and Organizations*, white paper, June 1998, http://www.microsoft.com/windows/platform/info/how2choose-mb.htm, November 1999.

[11] Microsoft Corp., *Microsoft Platform Software Development Kit Documentation*, 1999.

[12] Miller, B.P., Fredriksen, L. & So, B., "An empirical study of the reliability of Unix utilities," *Communications of the ACM*, 33(12): 32-43, December 1990.

[13] Miller, B.P., D. Koski, C. Pheow Lee, V. Maganty, R. Murthy, A. Natarajan & J. Steidl, *Fuzz Revisited: A Re-examination of the Reliability of UNIX Utilities and Services*, University of Wisconsin, CS-TR-95-1268, April 1995.

[14] Simon, R., *Windows NT Win32 API SuperBible*. Waite Group Press, Corte Modera, CA, 1997.

[15] Slabodkin, Gregory, "Software glitches leave Navy Smart Ship dead in the water," *Government Computer News*, http://www.gcn.com/archives/gcn/1998/july13/cov2.htm, July 13, 1998.

APPENDIX D

The Exokernel Operating System Architecture

by

Dawson R. Engler

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

October 1998

Author . . . . . .
Department of Electrical Engineering and Computer Science
May 18, 1998

Certified by . . . . . .
M. Frans Kaashoek
Associate Professor
Thesis Supervisor

Accepted by . . . . . .
Leonard A. Gould
Chairman, Departmental Committee on Graduate Students

The Exokernel Operating System Architecture

by
Dawson R. Engler

Submitted to the Department of Electrical Engineering and Computer Science
on May 18, 1998, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Computer Science and Engineering

Abstract

On traditional operating systems only trusted software such as privileged servers or the kernel can manage resources. This thesis proposes a new approach, the exokernel architecture, which makes resource management unprivileged but safe by separating management from protection: an exokernel protects resources, while untrusted application-level software manages them. As a result, in an exokernel system, untrusted software (e.g., library operating systems) can implement abstractions such as virtual memory, file systems, and networking.

The main thrusts of this thesis are: (1) how to build an exokernel system; (2) whether it is possible to build a real one; and (3) whether doing so a good idea. Our results, drawn from two exokernel systems [25, 48], show that the approach yields dramatic benefits. For example, Xok, an exokernel, runs a web server an order of magnitude faster than the closest equivalent on the same hardware, common unaltered Unix applications up to three times faster, and improves global system performance up to a factor of five.

The thesis also discusses some of the unusual techniques we have used to remove the overhead of protection. The most unusual technique, untrusted deterministic functions, enables an exokernel to verify that applications correctly track the resources they own, eliminating the need for it to do so. Additionally, the thesis reflects on the subtle issues in using downloaded code for extensibility and the sometimes painful lessons learned in building three exokernel-based systems.

Thesis Supervisor: M. Frans Kaashoek
Title: Associate Professor

The Exokernel Operating System Architecture

by
Dawson R. Engler

Submitted to the Department of Electrical Engineering and Computer Science
on May 18, 1998, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Computer Science and Engineering

Abstract

On traditional operating systems only trusted software such as privileged servers or the kernel can manage resources. This thesis proposes a new approach, the exokernel architecture, which makes resource management unprivileged but safe by separating management from protection: an exokernel protects resources, while untrusted application-level software manages them. As a result, in an exokernel system, untrusted software (e.g., library operating systems) can implement abstractions such as virtual memory, file systems, and networking.

The main thrusts of this thesis are: (1) how to build an exokernel system; (2) whether it is possible to build a real one; and (3) whether doing so a good idea. Our results, drawn from two exokernel systems [25, 48], show that the approach yields dramatic benefits. For example, Xok, an exokernel, runs a web server an order of magnitude faster than the closest equivalent on the same hardware, common unaltered Unix applications up to three times faster, and improves global system performance up to a factor of five.

The thesis also discusses some of the unusual techniques we have used to remove the overhead of protection. The most unusual technique, untrusted deterministic functions, enables an exokernel to verify that applications correctly track the resources they own, eliminating the need for it to do so. Additionally, the thesis reflects on the subtle issues in using downloaded code for extensibility and the sometimes painful lessons learned in building three exokernel-based systems.

Thesis Supervisor: M. Frans Kaashoek
Title: Associate Professor

“But I don't want to go among mad people,” Alice remarked.

“Oh, you can't help that,” said the Cat: “we're all mad here. I'm mad, you're mad.”

"How do you know I'm mad?" said Alice.

“You must be,” said the Cat, “or you wouldn't have come here.”

- Lewis Carroll (1832-1898), **Alice In Wonderland**

Acknowledgments

The exokernel project has been the work of many people. The basic principles of Chapter 2 and Aegis implementation come from a paper [25] written jointly with Frans Kaashoek and James O'Toole (which descended from my master's thesis, done under Kaashoek, with ideas initated by [26, 27]).

In contrast to Aegis, Xok has been written largely by others. Dave Mazieres implemented the initial Xok kernel. Thomas Pinckney Russell Hunt, Greg Ganger, Frans Kaashoek, and Hector Briceno further developed Xok and made ExOS into a real Unix system. Greg Ganger designed and implemented C-FFS [37] and the Cheetah webserver (based in part on [49]), the two linchpins of most of our application performance numbers. Ganger and Kaashoek oversaw countless modifications to the entire system. Eddie Kohler made an enormous contribution in our write up of these results. This work, described in [48], forms the basis for Chapter 5, and as a less primary source for Chapters 2— 4.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> And now that the legislators and the do-gooders have so futilely inflicted so many systems upon society, may they end up where they should have begun: may they reject all systems, and try liberty... — Frederic Bastiat
>
> It is hard to let old beliefs go. They are familiar. We are comfortable with them and have spent years building systems and developing habits that depend on them. Like a man who has worn eyeglasses so long that he forgets he has them on, we forget that the world looks to us the way it does because we have become used to seeing it that way through a particular set of lenses. Today, however, we need new lenses. And we need to throw the old ones away. — Kenich Ohmae

Traditional operating systems abstract and protect system resources. For example, they abstract physical memory in terms of virtual memory, disk blocks in terms of files, and exceptions and CPU in terms of processes. This organization has three significant benefits. First, it provides a portable interface to the underlying machine; applications need not care about the details of the underlying hardware. Second, it provides a large default functionality base, removing the need for application programmers to write device drivers or other low-level operating system code. Finally, it provides protection: because the operating system controls all application uses of resources, it can control application access to them, preventing buggy or malicious applications from compromising the system. Empirically, the ability to have multiple applications and users sharing the same machine is useful. Despite this organization's benefits, it has a serious problem: only privileged servers and the kernel can manage system resources. Untrusted applications are restricted to the interfaces and implementations of this privileged software. This organization is flawed because application demands vary widely. An interface designed to accommodate every application must anticipate all possible needs. The implementation of such an interface would need to resolve all tradeoffs and anticipate all ways the interface could be used. Experience suggests that such anticipation is infeasible and that the cost of mistakes is high [3, 9, 16, 25, 43, 79].

The *exokernel architecture* attacks this problem by giving untrusted application code as much safe control over resources as possible, thereby allowing orders of magnitude more programmers to innovate and use innovations, without compromising system integrity. It does so by dividing responsibilities differently from the way conventional systems do. Exokernels separate protection from management: they protect resources, applications manage them. An exokernel strives to move all functionality not required for protection out of privileged kernels and servers into unprivileged applications. For example, in the context of virtual memory, an exokernel protects physical pages and disk blocks used for paging, but defers the rest of management to applications (i.e., paging, allocation, fault handling, page table layout, etc.). The ideal exokernel makes untrusted software as powerful as a privileged operating system, without sacrificing either protection or efficiency.

Of course, not all applications need customized resource management. Instead of communicating with the exokernel directly, we expect most programs to be linked with libraries that hide low-level resources behind traditional operating system abstractions. However, unlike traditional implementations of these abstractions, library implementations are unprivileged and can therefore be modified or replaced at will. We refer to these unprivileged libraries as *library operating systems*, or libOSes. On the exokernels described in this thesis, libOSes implement virtual memory, file systems, networking, and processes. Applications are written on top of these libraries in a way similar to how they are written on top of current operating systems. As a result, applications can achieve appreciable speedups on an exokernel by simply linking against an optimized library operating system. Figure 1-1 illustrates a generic exokernel system.

clip= angle=270

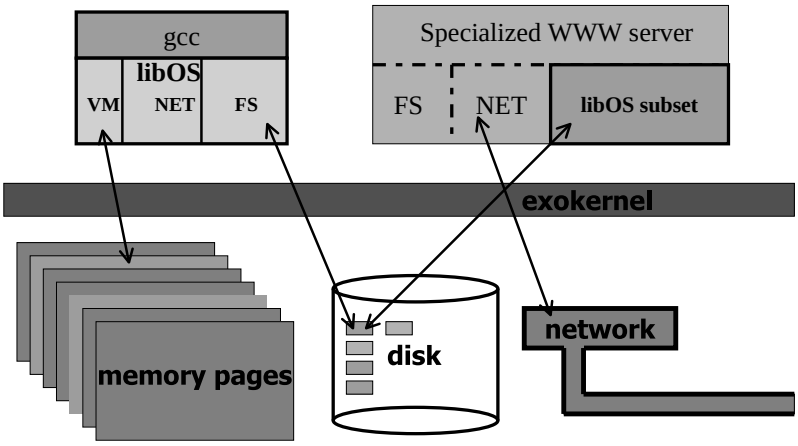

Figure 1-1: Possible exokernel system. Applications link against library operating systems (libOS), which provide standard operating system abstractions (virtual memory, files, network protocols, etc.). Because libOSes, are unprivileged, applications can also specialize them or write their own, as the web server in the picture has done. Because the exokernel provides protection, completely different libOSes can simultaneously run on the same system and safely share resources such as disk blocks and physical pages.

An exokernel retains the three benefits of traditional operating systems: default functionality and portability come from writing applications on top of a libOS, while protection comes from the exokernel, which guards all resources. In addition, we hope that the exokernel organization dramatically improves system innovation. We have four main reasons for this hope:

1. Fault-isolation: an error in a libOS only affects the applications using it, in contrast to errors in privileged operating systems, which can compromise the entire system. Thus, an exokernel significantly reduces the risk of using operating system innovations.

2. Co-existence: by design, multiple, possibly specialized, library operating systems can co-exist on the same exokernel system, in contrast to traditional systems, which by-and-large prevent more than one operating system running at a time. Thus, an exokernel enables innovation composition.

3. Increased implementor base: there are several orders of magnitude more systems programmers than privileged implementors (of oft-times proprietary operating systems). We hope the rate of innovation increases proportionally.

4. Increased user base: by making operating system software no different from other runtime libraries, the number of people with discretion to use an innovation increases by an even larger factor. Similarly, using innovations becomes a simple matter of linking in a new library rather than having to replace an entire system-wide operating system (and forcing all other users of the system to use it, and its bugs, in the process).

These four features remove many of the practical barriers facing operating system innovation development and deployment.

We do not assume that application programmers will modify operating system software as a matter of course. Instead, we regard libOS modification as similar to that of compilers: both are large, relatively complex pieces of software, not altered in the normal course of day-to-day programming. However, in the case of compilers, the enormous implementor and user community coupled with, unprivileged, fault-isolated compiler implementations has resulted in thousands of languages and implementations. For example, new languages such as Java, Tcl, Perl, and C++ have swept the implementor base every few years. Observed operating systems revolutions happen both on a more attenuated time scale, and with less dramatic scope. We hope that by making OS software more similar to compilers in the above ways that their evolution will become more similar as well.

An exokernel's success does not depend on a panoply of different operating systems. In our view, similar again to compilers and languages, there will be a few dominant operating systems but, importantly, the fact that they are unprivileged will enable them to evolve more readily than traditional systems.

1.1 Relation to other OS structures

There is a large literature on extensible operating systems, starting with the classic rationales by Lampson and Brinch Hansen [41, 53, 54]. Previous approaches to extensibility can be coarsely classified in to three groups: better microkernels, virtual machines, and downloading untrusted code into the kernel. We discuss each in turn and then relate exokernels to recent work in extensible operating systems.

The exokernel differs from traditional monolithic systems in that it places the bulk of operating system functionality in unprivileged libraries. It similarly differs from a microkernel in that, while both organizations move code out of the kernel, a microkernel pushes code into privileged servers, which applications cannot modify. In some sense, the principal goal of an exokernel—giving applications control—is orthogonal to the question of monolithic versus microkernel organization. If applications are restricted to inadequate interfaces, it makes little difference whether the implementations reside in the kernel or privileged user-level servers [39, 39]; in both cases applications lack control. For example, it is difficult to change the buffer management policy of a shared file server. In many ways, servers can be viewed as fixed kernel subsystems that happen to run in user space. Whether monolithic or microkernel-based, the goal of an exokernel system remains for privileged software to provide interfaces that do not limit the ability of unprivileged applications to manage their own resources.

Hydra was the most ambitious early system to have the separation of kernel policy and mechanism as one of its central tenets [94]. An exokernel takes the elimination of policy one step further by removing “mechanism” wherever possible. This process is motivated by the insight that mechanism *is* policy, albeit with one less layer of indirection. For instance, a page-table is a very detailed policy that controls how to translate, store and delete mappings and what actions to take on invalid addresses and accesses. [1]

Some newer microkernels push the kernel interface closer to the hardware [16, 42, 73], obtaining better performance and robustness than previous microkernels and allowing for a greater degree of flexibility, since shared monolithic servers can be broken into several servers. Techniques to reduce the cost of shared servers by improving IPC performance, moving code from servers into libraries, mapping read-only shared data structures, and batching system calls [8, 39, 58, 61] can also be successfully applied in an exokernel system.

Virtual machines [12, 34, 38] (VMs) are an OS structure in which a privileged virtual machine monitor (VMM) isolates less privileged software in emulated copies of the underlying hardware. Virtual machines and exokernels differ in two main ways. First, virtual machines emulate, exokernels do not. Emulation hides information. This can lead to ineffective use of hardware resources; for instance, the VMM has no way of knowing if a VM no longer needs a particular virtual page. In contrast, an exokernel attempts to expose all information about the system and explicitly communicates with the library operating system rather than presenting a virtual facade and making its own resource management decisions. Second, VMs can only share resources through remote communication protocols. This prevents VMs from sharing many OS abstractions such as processes or file descriptors with each other. Thus, VMMs confine specialized operating systems and associated processes to isolated virtual machines. Rather than partitioning the machine into disjoint pieces, an exokernel allows non-trusting applications to use customized libOSes

[1] The OS community has no rigorous definition of either “policy” or “mechanism.” In its most general sense, policy refers to the goals of a computer system (e.g., what security threats to resist, what resources to protect). In context of extensible operating systems, policy refers to the algorithm used to make security or resource management decisions, while mechanism refers to the machinery used to implement a particular policy. For example, a virtual memory paging policy is to evict least recently used pages to disk. The data structures used to do so, such as page tables and a sorted page list, would be mechanism. Similar to the concepts of code and data, there is no clear fundamental difference between these two notions.

to share resources (such as physical memory and disk blocks) without sacrificing a single view of the machine,

Downloading code into the kernel is another approach to extensibility. In many systems only trusted users can download code, either through dynamically-loaded kernel extensions or static configuration [33, 43]. In the SPIN and Vino systems, any user can safely download code into the kernel [9, 78]. Safe downloading of code through type-safety [9, 75] and software fault-isolation [78, 89] is complementary to the exokernel approach of separating protection from management. Exokernels use downloading of code to let the kernel leave decisions to untrusted software [25].

In addition to these structural approaches, much work has been done on better OS abstractions that give more control to applications, such as user-level networking [88, 82], lottery scheduling [90], application-controlled virtual memory [44, 55] and file systems [13, 72]. All of this work is directly applicable to libOSes.

1.1.1 Recent extensible operating systems

SPACE is a "submicro-kernel" that provides only low-level kernel abstractions defined by the trap and architecture interface [73]. Its close coupling to the architecture makes it similar in many ways to an exokernel, but we have not been able to make detailed comparisons because its design methodology and performance have not yet been published.

The SPIN project is building a microkernel system that allows applications to make policy decisions [9] by safely downloading *extensions* into the kernel. Unlike SPIN, the focus in the exokernel architecture is to obtain flexibility and performance by securely exposing low-level hardware primitives rather than extending a traditional operating system in a secure way. As a result, exokernel interfaces tend to be lower level and, thus, grant more control to applications.

Anderson [3] makes a clear argument for application-specific library operating systems and proposes that the kernel concentrate solely on the adjudication of hardware resources. The exokernel design addresses how to provide secure multiplexing of physical resources in such a system, and moves the kernel interface to a lower level of abstraction. In addition, our exokernel systems demonstrate that low-level secure multiplexing and library operating systems can offer excellent performance.

Like an exokernel, the Cache Kernel [16] provides a low-level kernel that can support multiple application-level operating systems. To the best of our knowledge the Cache Kernel and ExOS, the libOS used on three generations of exokernels [25, 48], are the first general-purpose library operating systems implemented in a multiprogramming environment. The difference between the Cache Kernel and an exokernel is mainly one of high-level philosophy. The Cache Kernel focuses primarily on reliability, rather than securely exporting hardware resources to applications. As result, it is biased towards a server-based system structure. In our experience, servers become de facto privileged in that their functionality cannot be overridden by applications.

The Nemesis kernel [76, 56] has many similarities to an exokernel, despite a vast difference in goals. Nemesis is designed to improve quality-of-service for multimedia applications. The problem it attacks is that traditional systems make resource accounting difficult in that code running on behalf of an application may be strewn throughout a collection of servers and the kernel. Like an exokernel, their solution is to push operating system functionality into applications. In this way, code running on behalf of the application typically runs within the application itself, which then can trivially be charged for its resource consumption. The primary parallels between Nemesis and an exokernel system are low-level interfaces for resources needed by multimedia applications (events, network, disk, and CPU), and a reliance on library operating systems. However, the difference in goals leads to stark contrasts. While Nemesis is designed to simplify accounting, an exokernel aims to improve innovation. It does so by ceding *all* control not needed for protection to applications. Thus, an exokernel's interfaces grant more pervasive power to applications and also promote sharing of resources (so that different implementations can safely share the same state). For example, applications have only limited control over virtual memory, since Nemesis forces a single-address space on the entire system. Similarly, Nemesis disk multiplexing lacks the power of our disk subsystem, XN (discussed in Chapter 4). We know of no experimental results for Nemesis, which prevents us from performing a more detailed comparison of the relative strengths of the two approaches.

1.2 The focusing questions of this thesis

An exokernel attempts to make unprivileged library operating systems as powerful as privileged operating systems. The main thrusts of this thesis are:

1. How to build an exokernel system.

2. Whether it is possible to build a real one.

3. Whether doing so is a good idea.

The first half of the thesis, Chapters 2 and 4, focuses on how to build an exokernel system, and the later chapters on the approach's efficacy.

The text below provides an overview of the questions each issue raises, along with a sketch of their associated answers.

1.2.1 How to build an exokernel?

Traditionally, operating systems have provided a high-level interface to unprivileged software. Chapter 2 provides a constructive methodology for how to lower this interface to a level sufficient for libOSes to build abstractions such as networking, virtual memory and file systems. The six principles of this methodology are: (1) separate management from protection; (2) expose all hardware to applications; (3) expose resource allocation, placing it under the discretion of applications; (4) expose revocation, thereby letting applications determine what resources to relinquish; (5) protect at a fine-grained level, making sharing and management lightweight; and (6) expose information, such as hardware capabilities, physical names, and global system statistics.

A key aspect of exokernel design is leaving implementation decisions to the client. Rather than focusing design on deciding how to implement a policy or mechanism, focus on constructing an interface that leaves all interesting decisions to applications. An interface that merely checks that an application has performed an operation correctly frequently gives greater freedom for important decisions, as well as being simpler to implement. In some sense, exokernel design focuses on the art of transmuting the imperative (deciding how to implement a policy or mechanism in the kernel) to the declarative (specifying *what* interface an application must implement, and checking that it does so correctly).

Chapter 3 contains a number of examples of how to apply this methodology, and Chapter 4 discusses an extended example of how to multiplex a hardware disk, which has been the most challenging resource for us to handle.

The following four subsections discuss important subproblems of exokernel construction.

How to recapture resource semantics?

By dislodging OS code into libraries, an exokernel also ejects a significant portion of the code that understands resource semantics. For example, in the context of networking, because the exokernel dislocates network protocol code (e.g., TCP/IP and UDP) into untrusted libraries, it no longer understands packet semantics. As a result, it lacks the information necessary to decide which application owns what message: a decision it must make if it is to implement protection.

We have used downloaded code as a powerful mechanism to allow untrusted software to convey these semantics back into the exokernel in a general way. In the context of networking we use packet filters (see Chapter 2), for disk, deterministic meta data interpreters (see Chapter 4).

From a different perspective, an exokernel, by “uploading” code into the application, removes the need to protect many pieces of state, and, thus, trivially need not understand their semantics.

Chapter 6 reflects on our experience using downloading code: when downloaded code was superior (or inferior) to a functional interface, our mistakes, and a number of surprises, visible only in much delayed hindsight. This technique has subtle implications, resulting from, among other things: the asymmetry in trust between the extension and the host; the fact that code for most useful languages is Turing complete, while most procedural interfaces are decidedly not; and that downloaded code, because it can be restricted, can be granted abilities that unrestricted external application code cannot. A consequence of this latter attribute is that most of our uses of downloaded code have little to do with speed but rather with granting applications power otherwise not possible.

How to protect shared state?

Traditional operating systems encapsulate and enforce well-formed updates to state shared amongst processes. For example, a Unix kernel performs modifications on shared file descriptors, the file name cache, the buffer cache, etc. In contrast, an exokernel dislocates such state into unprivileged applications, which then modify it, potentially incorrectly. How then can invariants on shared state be enforced? We have used a plethora of techniques to enforce this.

The most general solution is to apply the exokernel precepts recursively: divide libraries into privileged and unprivileged parts, where the "privileged" part contains all code required for protection, and must be used by all applications using a piece of state, while the unprivileged contains all management code and can be replaced by anyone. Forcing applications to use the privileged portion of a libOS can be done by placing it in a server, using a restricted language and trusted compiler, or downloading code into the kernel. We have used all three.

In the more common case, library operating systems can frequently be designed for localization of state and to use standard fault-isolation techniques (such as type-safe languages and memory protection).

In many cases, the desire to protect shared state with mechanisms more elaborate and read and write access checks makes little technical sense: if the raw data itself can be corrupted, it is unlikely that there is any reason to preserve high-level invariants on the resultant garbage. One of the most common situations where this dynamic arises is in the protection of shared bookkeeping data structures. Consider the case of a shared buffer (say a Unix pipe), that uses a buffer record to track the number of bytes in the buffer. At one level, one would like to guarantee that an application decrements the byte count on data removal, and increments it on insertion. However, while this desire is sensible from a software engineering perspective, it is not a protection issue. Since there are no guarantees on the sensibleness of the inserted data, knowing how many bytes of garbage are in the buffer gains no security. In practice, social, rather than technical, methods guarantee these sort of invariants — e.g., assemblers adhere to standard object code formats rather than going through a set of system enforced methods for laying out debugging and linking information.

How to protect without overhead?

An exokernel adds another layer to the system in that it takes operating system code, which formerly ran on bare hardware, and places it in an unprivileged library, which runs on top of an exokernel which runs on bare hardware. Since performance motivates exokernels, the cost of this extra layer must be negligible.

Fortunately, for most resources, this overhead has been a non-issue. Exokernel implementation of virtual memory, networking, exceptions, and process management have all performed well, without aggressive tuning. In fact, on the first exokernel system we built, Aegis [25], if an exokernel primitive did not perform significantly faster than that of a traditional system from the beginning, we looked for "what was wrong." [2]

In practice, protection does not impose overhead. As we discuss in Chapter 5, it tends to be off of the critical path, coupled to expensive operations that dwarf its overhead, or recovered by other features of the architecture (e.g., libOSes make many system calls into function calls).

Can applications be trusted to track ownership?

At its most basic level, protection requires a table lookup to map a principal to access rights. While maintaining this table adds little overhead for most resources, for others, such as a disk, it is infeasible. Chapter 4 presents the reasons for this in detail.

This problem's solution comes from noticing that correct applications track what resources they have access to, rendering operating system bookkeeping redundant. For example, a library file system tracks the disk blocks each file maps to. Thus, if the operating system can reuse the application's data structures, it can eliminate this redundancy. In the case of file systems, if the exokernel can reuse the library file system's "meta data" (the data structures it uses to map files to blocks), then it need not perform duplicate bookkeeping. We have invented an online verification technique that lets an exokernel verify that library file systems correctly track what disk blocks they own, without the operating system having to understand how they do so.

1.2.2 Can you build a real exokernel system?

Does the exokernel organization only work for relatively simple, isolated resources such as physical memory, or can it apply to more difficult shared resources such as disk? Can one build an exokernel system? Fortunately for this thesis, one can. There have been three exokernel systems built so far, in increasing verisimilitude: Aegis [25], Glaze [60], and Xok [48]. Most of our performance data and examples come from Xok and ExOS, its default libOS. While ExOS does not handle some Unix corner cases, it is not a toy either. For example, it runs most Unix applications (e.g., perl,

[2] Aegis was roughly a factor of ten faster than a mature monolithic system. This difference held for exceptions, virtual memory primitive operations, inter-process communication, system calls, and even operations with large fixed costs, such as message round trip times over a 10Mb/s Ethernet. [25]

gcc, telnet, and csh) without modification. This fact can be seen in that the bulk of our performance data comes from application end-to-end numbers rather than micro-benchmarks.

1.2.3 Are exokernels a good idea?

Chapter 5 (and to a lesser degree, Chapter 3) focuses primarily on evaluating exokernel efficacy. We use a performance driven approach; our experiments address four questions, discussed below.

Do common, unaltered applications benefit?

If an exokernel only improves the performance of strange niche applications or requires that applications be modified to benefit, then its usefulness is severely diminished.

Our experiments show that an exokernel matters, even for common applications. We measure the performance of a software development workload made up of mainstream applications (most are found in "/usr/bin" on a Unix system). When linked against an optimized library file system and run on our exokernel system, these programs run up to three times faster than identical versions executed on competitive monolithic operating systems.

In general, normal applications benefit from an exokernel by simply linking in an optimized libOS that implements a standard interface. Thus, even without modifications, they still profit from an exokernel. [3] While an exokernel allows experimentation with completely different OS interfaces, a more important result may be improving the rate of innovation of implementations of existent interfaces.

Does a libOS run slower than an OS?

An exokernel provides extreme flexibility. Rightfully, any systems builder would expect that this flexibility would have a performance cost. In particular, given the same application code and same OS code (placed in a libOS on an exokernel, in the kernel on a monolithic system), does a traditional system provide superior performance? Or, phrased another way, if an optimization is done on an exokernel and gives a factor of four, would the same optimization done on a traditional OS give even more?

To partially answer this question we have taken the library file system discussed in the previous experiment and re-implemented it in the kernel of a traditional, monolithic operating system. We then run the same workload on it. Our results show that, at least in this case, an exokernel does not pay for its structure. Or, in the alternative way, that an optimization done on an exokernel can give the same performance improvement as done on a traditional system.

Are aggressive applications ten times faster?

In part, the exokernel architecture was motivated by the ability to perform application-specific or, more commonly, domain-specific optimizations. [4] Two natural questions, then, are first, does an exokernel give sufficient power that interesting optimizations can be done? Second, do domain-specific optimizations yield significant improvements or do they only give "noise" level improvements?

Experiments show that an exokernel enables domain-specific optimizations that give order-of-magnitude performance improvements [48]. Our specific results come from an interface designed for fast I/O, which improves the performance of applications such as web servers.

Does local control lead to bad global performance?

An exokernel gives applications significantly more control than traditional operating systems do. A positive aspect of this power transfer is that applications can perform optimizations not previously possible. A negative aspect is that their selfish optimization efforts could destroy system performance. Thus a key question is: can an exokernel reconcile local control with good global performance?

It can, for two main reasons. First, when an exokernel structure enables improved application performance there will be more resources to go around, and thus, global performance will improve. Second, an exokernel mediates allocation and revocation of resources. Therefore it has the power to enforce any global policy that a traditional

[3] In principle, given sufficient resources a traditional OS can implement any innovation done on an exokernel. However, in practice, we expect an exokernel system to evolve significantly faster than traditional systems.

[4] In fact, the original exokernel paper [25] focuses almost exclusively on application-specific optimizations rather than improving the rate of general-purpose innovations, which we have come to regard as a more significant benefit.

operating system can. The single new challenge it faces is deriving information lost by dislocating abstractions into application space. For example, traditional operating systems manage both virtual memory and file caching. As a result, they can perform global resource management of pages that takes into account the manner in which a page is being used. In contrast, if an exokernel dislocates virtual memory and file buffer management into library operating systems it no longer can make such distinctions. While such information matters in theory, in practice we have found it either unnecessary or crude enough that no special methods have been necessary to derive it. However, whether this happy situation always holds is an open question.

More concretely, our experiments demonstrate that: (1) given the same workload, an exokernel performs comparably to widely used monolithic systems, and (2) that when local optimizations are performed, that whole system performance improves, and can do so significantly.

1.2.4 Summary

Taken as a whole, these experiments show that, compared to a traditional system, an exokernel provides comparable to significantly superior performance while simultaneously giving tremendous flexibility.

1.3 Concerns

Below, we discuss five concerns about the exokernel architecture: (1) that an it will compromise portability, (2) that it will lead to a Babel of incompatible libOSes, (3) that libOSes are too costly to specialize, (4) that they consume too much space per application, and (5) that the modification of a “free software” OS provides a “good enough” way to innovate.

First, it is a virtue for an exokernel to expose the details of the system's hardware and software. Our exokernels expose information such as the number of network and scatter/gather DMA buffers available and TLB entry format. Naively handled, this exposure can compromise portability. Fortunately, traditional techniques for retaining portability work just as well on exokernel systems. While applications that use an exokernel interface directly will not be portable, those that use library operating systems that implement standard interfaces (*e.g.*, POSIX) are portable across any system that provides the same interface. Library operating systems themselves can be made portable by designing them to use a low-level machine-independent layer to hide hardware details. This technique has been widely used in the context of more traditional operating systems [74, 21].

As in traditional systems, an exokernel can provide backward compatibility in three ways: one, binary emulation of the operating system and its programs; two, implementing its hardware abstraction layer on top of an exokernel; and three, re-implementing the operating system's abstractions on top of an exokernel.

Second, since library operating systems are unprivileged, anyone can write one. An obvious problem is what happens to system coherence under an onslaught of “hand rolled” operating systems? Similar to the previous challenge, applications already solve the problem of chaos from freedom through the use of standards and good programming practices. The advantage of standards on paper rather than hard wired in the kernel is that they can be re-implemented and specialized, selected amongst, or, in the case of the worst ones, ignored.

Third, does specialization require writing an entire library operating system? Our experience shows that it does not. All of the application performance improvements in this thesis, even the most dramatic, were achieved by inserting the specialized subsystem within the default libOS, rather than rewriting an entire libOS from scratch. Given a modular design, other library operating systems should allow similar specialization. It is possible that object-oriented programming methods, overloading, and inheritance can provide useful operating system service implementations that can be easily specialized and extended, as in the VM++ library [51].

Most libOS code is no more difficult to modify than other systems software such as memory allocators. It tends to be conceptually shallow, and concerned chiefly with managing various mappings (page tables, meta data, file tables, etc.). As such, once operating system software is removed from the harsh environ of the kernel proper, modification does not require extraordinary competence.

Fourth, if each application has its own library operating system linked into it, what about space? As in other domains with large libraries, exokernel systems can use dynamic linking and shared libraries to reduce space overhead. In our experience, these techniques reduce overheads to negligible levels. For example, our global performance experiments place a severe strain on system memory resources, yet an exokernel is comparable to or surpasses the conventional systems we compare against.

Finally, what about Linux, FreeBSD, and the rest of the public-domain operating systems? Do they not evolve as quickly as an exokernel will? We believe that an exokernel structure has important software engineering benefits over these systems. First, libraries are fault-isolated from each other, allowing different operating systems to co-exist on the same system, something not possible with the current crop of operating systems. Second, library development is much much easier than kernel hacking. For example, standard debugging tools work, and errors crash only the application using the library instead of the system, allowing easy post-mortem examination. Third, all users have the power to use an innovation, simply by linking an application against a libOS.

1.4 Summary

The exokernel in a nutshell:

1. What: anyone can safely manage resources.
2. How: separate protection from management. An exokernel protects resources, applications manage them. Everything not required for protection is moved out of privileged kernels and servers into untrusted application libraries. This is the one idea to remember for an exokernel, all other features are details derived from it.

   The exokernel ideal: the libOS made as powerful as privileged OS, without sacrificing performance or protection.
3. Why? Innovation. An exokernel makes operating systems software co-exist, unprivileged, and modifiable by orders of magnitude more programmers.

The following chapter articulates the principles of the exokernel methodology. Chapter 3 draws on examples from two exokernel systems, Aegis and Xok, to show how these principles apply in practice. The most challenging problem for protection, disk multiplexing, is discussed in detail in Chapter 4. Chapter 5 presents application performance results that show that that such separation is profitable. We reflect on using downloaded code for extensibility in Chapter 6. Finally, Chapter 7 discusses future work, reflects on general lessons learned while building exokernel systems, and concludes.

Chapter 2

How to build an exokernel: principles

> We all live in the protection of certain cowardices which we call our principles . — Mark Twain (1835-1910)
>
> It is impossible to design a system so perfect that no one needs to be good. — T. S. Eliot:

The goal of an exokernel is to give as much safe, efficient control over system resources as possible. The challenge for an exokernel is to give library operating systems maximum freedom in managing resources while protecting them from each other; a programming error in one library operating system should not affect another library operating system. To achieve this goal, an exokernel separates protection, which must be privileged, from management, which should be unprivileged.

Figure 2-1 shows a simplified exokernel system that is running two applications: an unmodified UNIX application linked against the ExOS libOS and a specialized exokernel application using its own TCP and file system libraries. Applications communicate with the kernel using low-level physical names (e.g., block numbers); the kernel interface is as close to the hardware as possible. LibOSes handle higher-level names (e.g., file descriptors) and supply abstractions. Protection in this figure consists of forcing applications to only read or write (cached) disk blocks that they have access rights to. The kernel does this by guarding all operations on both the disk blocks themselves and on the physical pages that hold cached copies.

In general, resource protection consists of three tasks: (1) tracking ownership of resources, (2) ensuring protection by guarding all resource usage or binding points, and (3) revoking access to resources. To prevent rogue applications from destroying the system, most exokernels protect physical resources such as physical memory, the CPU, and network devices (to prevent theft of received messages and corruption of not-yet-transmitted ones). Additionally, they must guard operations that change privileged state such as the ability to execute privileged instructions, and writes to "special" memory locations that control devices. Finally, they may guard more abstract resources such as bandwidth.

To make this discussion concrete, consider the protection of physical pages. The core requirement is that the exokernel must check that every read and write to any page has sufficient access rights. On modern architectures, encapsulation involves forcing all memory operations to either go through the kernel (which checks permissions manually) or through translation hardware (the TLB), which checks them against a set of cached pages. To ensure that malicious processes do not corrupt the TLB, the kernel must check modifications of it as well. In practice this means that applications cannot issue privileged instructions directly, but must go through the kernel. The next three chapters give many examples of resources to protect and ways to do so.

Applications cannot override privileged software. This inflexibility makes protection possible. However, if it spills over into non-protection areas, applications will needlessly lose the ability to control important decisions. Thus, an exokernel only overrides application decisions to the extent necessary to build a secure system.

This chapter describes five principles that give this declarative goal operational teeth. These principles illustrate the mechanics of exokernel systems and provide important motivation for many design decisions discussed later in this thesis. To better understand what the architecture is, we then show what it is not. The remainder of the chapter discusses how an exokernel protects the state of higher-level abstractions, and performs resource revocation, and then presents a general technique, *secure bindings*, we have used to make protection efficient.

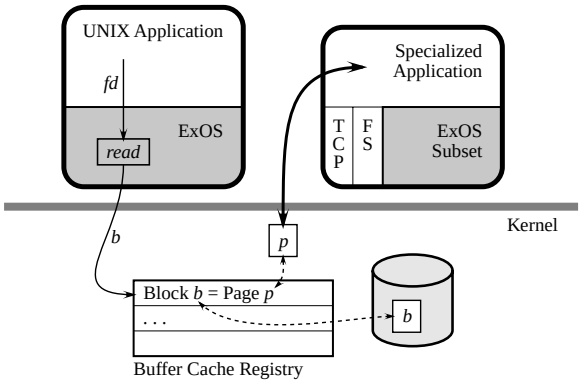

Figure 2-1: A simplified exokernel system with two applications, each linked with its own libOS and sharing pages through a buffer cache registry.

2.1 Exokernel principles

The goal of an exokernel is to give efficient control of resources to untrusted applications in a secure, multi-user system. We follow these principles to achieve this goal:

Separate protection and management. Exokernels restrict resource management to functions necessary for protection: allocation, revocation, sharing, and the tracking of ownership. Giving applications control over all non-protection mechanisms and policies makes the system "optimally" extensible, in that any further application customization would compromise system integrity.

In general, an exokernel strives to provide applications access to all operations that a privileged OS has. This requires providing ways for libOSes to recapture aspects of the kernel's privileged execution context, such as being tightly coupled to hardware interrupts, which they have lost by running as application software. (From another perspective, the privileged exokernel adds a layer of indirection between a libOS and hardware events. The exokernel must take steps to ensure that this layer does not preclude flexible application control of hardware.)

Expose hardware. Exokernels give applications protected access to all resources. Applications have access to privileged instructions, hardware DMA capabilities, and machine resources. The resources exported are those provided by the underlying hardware: physical memory, the CPU, disk memory, translation look-aside buffer (TLB), and addressing context identifiers. Applications have full view of resource attributes, such as the number, format, and current set of TLB mappings or the number and size of incoming and outgoing network buffers, as well as the ability to modify them (inserting TLB entries). This principle extends to less tangible machine resources such as interrupts, exceptions, and cross-domain calls.

An exokernel exposes resources and events at the lowest possible level required for protection—ideally, at the level of hardware (disk blocks, physical pages, etc.) rather than rather than encapsulating them in high-level abstractions such as files, or Unix signal or RPC semantics.

The following two principles fall from a general pattern: involve applications in important decisions rather than subjecting them to a hard wired policy.

Expose allocation. Applications allocate resources explicitly. The kernel allows specific resources to be requested during allocation, giving applications control over abstraction semantics and performance decisions, such as when and which disk blocks to allocate.

Expose revocation. Exokernels expose revocation policies to applications. They let applications choose which instance of a resource to give up. Each application has control over its set of physical resources. Exokernels allow applications to revoke resources from processes they control, thereby allowing enforcement of local policies.

Protect fine-grained units. To make resource management and sharing lightweight, exokernel protection is fine-grained (e.g., at the level of disk blocks rather than partitions). Fine-grained protection reduces the overhead of mapping high-level abstractions down to the discrete resource units that an exokernel protects (e.g., mapping files to disk blocks, virtual memory to pages, network messages to message buffers, etc.). Coarse-grained protection wastes space as well as time. For example, in the context of disk, by increasing the length of disk arm sinks after increasing

internal fragmentation through coarse-grained allocation. Coarse-grained protection also makes sharing clumsy, since resources cannot be shared at a finer-granularity than the exokernel protects.

Expose names. Exokernels use physical names wherever possible. Physical names capture useful information and do not require potentially costly or race-prone translations from virtual names. For example, physical disk blocks encode position, one of the most important variables for file system performance. Virtual names, in contrast, would not necessarily encode this information and would require on-disk translation tables. Using these tables increases disk accesses (extra writes for persistence, extra reads for name translation) and protecting them across reboots degrades performance (e.g., by ordering writes to disk, and causing multiple writes).

In practice, physical names also provide a uniform mechanism for optimistic synchronization. They allow locks to be replaced with checks that "expected" conditions hold (that a directory name maps to a specific disk block and has not been deleted or moved, etc.).

Expose information. Exokernels expose system information, and collect data that applications cannot easily derive locally. For example, applications can determine how many hardware network buffers there are or which pages cache file blocks. An exokernel might also record an approximate least-recently-used ordering of all physical pages, something individual applications cannot do without global information. Additionally, exokernels export book-keeping data structures such as free lists, disk arm positions, and cached TLB entries so that applications can tailor their allocation requests to available resources.

We have found it useful in practice to provide space for protected application-specific data in kernel data structures. For example, the structures describing an execution environment are improved if application-specific data can be associated with them (e.g., to record the numeric id of an process' parent, its running status, program name, etc.).

These principles apply not just to the kernel, but to any component of an exokernel system. Privileged servers should provide an interface boiled down to just what is required for protection.

2.1.1 Policy

An exokernel hands over resource policy decisions to library operating systems. Using this control over resources, an application or collection of cooperating applications can make decisions about how best to use these resources. However, as in all systems, an exokernel must include policy to arbitrate between competing library operating systems: it must determine the absolute importance of different applications, their share of resources, *etc.* This situation is no different than in traditional kernels. Appropriate mechanisms are determined more by the environment than by the operating system architecture. For instance, while an exokernel cedes management of resources to library operating systems, it controls the allocation and revocation of these resources. By deciding which allocation requests to grant and from which applications to revoke resources, an exokernel can enforce traditional partitioning strategies, such as quotas or reservation schemes. Since policy conflicts boil down to resource allocation decisions (*e.g.*, allocation of seek time, physical memory, or disk blocks), an exokernel handles them in a similar manner.

2.2 Kernel support for protected abstractions

Many of the resources protected by traditional operating systems are themselves high-level abstractions. Files, for instance, consist of meta data, disk blocks, and buffer cache pages, all of which are guarded by access control on high-level file objects. While exokernels allow direct access to low-level resources, exokernel systems must be able to provide UNIX-like protection, including access control on high-level objects where required for security. One of the main challenges in designing exokernels is to find kernel interfaces that allow such higher-level access control without either mandating a particular implementation or hindering application control of hardware resources.

We have met this challenge by providing both protection machinery in the kernel and ways for applications to download code to augment this machinery. In the former category, our exokernels provide provides software abstractions to bind hardware resources together. For example, as shown in Figure 2-1, the Xok buffer cache registry binds disk blocks to the memory pages caching them. Applications have control over physical pages and disk I/O, but can also safely use each other's cached pages. Our exokernel's protection mechanism guarantees that a process can only access a cache page if it has the same level of access to the corresponding disk block.

More generally, by allowing applications to download code, an exokernel gives them a mechanism to encapsulates an abstraction's state and can thereby enforce invariants on it. This strategy is required for abstractions whose protection does not map to hardware abstractions. For example, files may require valid updates to their modification times. Our

oldest use of downloaded code is packet filters. Packet filters are application-written code fragments that applications download into the kernel to select what incoming network packets they want.

However, while these software abstractions reside in the kernel, they could also be implemented in trusted user-level servers. This microkernel organization would cost many additional (potentially expensive) context switches. Furthermore, partitioning functionality in user-level servers tends to be more complex.

From one perspective, downloaded code provides a conduit for pushing semantics across the user-kernel trust boundary. This activity is more important for exokernels than traditional systems. By dislodging OS code into libraries, an exokernel also removes the code that understands resource semantics and, thus, lose the ability to protect these resources. For networking, because an exokernel ejects the code that understands network protocol semantics it no longer understands how to bind incoming packets to applications. Packet filters provide a way to recapture these semantics by encapsulating them in a black box (the filter), which is then injected into the kernel. The exokernel thus trades an intractable problem — anticipating all possible invariants that must be enforced — for one that is tractable: testing the downloaded code for trustworthiness (in this case, that one filter will not steal another's packets). Chapter 6 discusses this perspective in more detail.

The key to these exokernel software abstractions is that they neither hinder low-level access to hardware resources nor unduly restrict the semantics of the protected abstractions they enable. Given these properties, a kernel software abstraction does not violate the exokernel principles.

Protected sharing

The low-level exokernel interface gives libOSes enough hardware control to implement all traditional operating system abstractions. Library implementations of abstractions have the advantage that they can trust the applications they link with and need not defend against malicious use. The flip side, however, is that a libOS cannot necessarily trust all other libOSes with access to a particular resource. When libOSes guarantee invariants about their abstractions, they must be aware of exactly which resources are involved, what other processes have access to those resources, and what level of trust they place in those other processes.

As an example, consider the semantics of the UNIX fork system call. It spawns a new process initially identical to the currently running one. This involves copying the entire virtual address space of the parent process, a task operating systems typically perform lazily through copy-on-write to avoid unnecessary page copies. While copy-on-write can always be done in a trusted, in-kernel virtual memory system, a libOS must exercise care to avoid compromising the semantics of fork when sharing pages with potentially untrusted processes. This section details some of the approaches we have used to allow a libOS to maintain invariants when sharing resources with other libOSes.

Our latest exokernel, Xok, provides three mechanisms libOSes can use to maintain invariants in shared abstractions. First, *software regions*, areas of memory that can only be read or written through system calls, provide sub-page protection and fault isolation. Second, Xok allows on the-fly-creation of *hierarchically-named capabilities* and requires that these capabilities be specified explicitly on each system call [62]. Thus, a buggy child process accidentally requesting write access to a page or software region of its parent will likely provide the wrong capability and be denied permission. Third, Xok provides robust critical sections: inexpensive critical sections that are implemented by logically disabling software interrupts [10]. Using critical sections instead of locks eliminates the need to trust other processes.

Three levels of trust determine what optimizations can be used by the implementation of a shared abstraction.

Optimize for the common case: Mutual trust. It is often the case that applications sharing resources place a considerable amount of trust in each other. For instance, any two UNIX programs run by the same user can arbitrarily modify each others' memory through the debugger system call, ptrace. When two exokernel processes can write each others' memory, their libOSes can clearly trust each other not to be malicious. This reduces the problem of guaranteeing invariants from one of security to one of fault-isolation, and consequently allows libOS code to resemble that of monolithic kernels implementing the same abstraction.

Unidirectional trust. Another common scenario occurs when two processes share resources and one trusts the other, but the trust is not mutual. Network servers often follow this organization: a privileged process accepts network connections, forks, and then drops privileges to perform actions on behalf of a particular user. Many abstractions implemented for mutual trust can also function under unidirectional trust with only slight modification. In the example of copy-on-write, for instance, the trusted parent process must retain exclusive control of shared pages and its own page tables, preventing a child from child making copied pages writable in the parent. While this requires more page faults in the parent, it does not increase the number of page copies or seriously complicate the code.

Mutual distrust. Finally, there are situations where mutually distrustful processes must share high-level abstractions with each other. For instance, two unrelated processes may wish to communicate over a UNIX domain socket, and neither may have any trust in the other. For OS abstractions that can be shared by mutually distrustful processes, the most general solution is to apply the exokernel precepts recursively: divide libraries into privileged and unprivileged parts, where the "privileged" part contains all code required for protection, and must be used by all applications using a piece of state, while the unprivileged contains all management code and can be replaced by anyone. Forcing applications to use the privileged portion of a libOS can be done by placing it in a server, using compiler techniques, or downloading code it into the kernel. We have used all three.

Fortunately, sharing with mutual distrust occurs very infrequently for many abstractions. Many types of sharing occur only between child and parent processes, where mutual or unidirectional trust almost always holds.

2.3 Implementation-defined Decisions

The exokernel architecture defines a family of implementations. To understand it, it helps to understand what it is not. It leaves the following five decisions to the system designer.

What to protect. While it is hard to imagine a usable system that elides memory and disk protection, the exokernel architecture does not mandate what should be protected. It only says that once this decision has been made, that all functionality not needed for protection should be implemented by untrusted software. Example resources that our exokernel implementations deliberately do not protect are: bandwidth quality of service guarantees, deadline guarantees, or any number of restrictions required by real time systems.

What global system policies to implement. Every system includes global policies enforced on all applications that decide which processes to revoke resources from and which requests to grant or deny (even if this policy is no policy at all). What specific global policy is enforced — whether it be optimized for interactive performance, throughput, real time, etc. — is orthogonal to the architecture. However, given a policy, an exokernel designer strives to involves application decisions in its implementation.

How to protect. While we provide examples of how our exokernel systems have multiplexed memory, disk, network, etc., these ways are not the only ones. Again, what is important in an exokernel is allowing untrusted software to implement everything not needed for protection; the designer has discretion in how he achieves this. This decision includes how to track access rights: our first exokernel used self-authenticating capabilities [15], then access control bitmaps [16], our second hierarchical capabilities [63].

The level of protection. An exokernel interface may contain high-level primitives. The exokernel principles are guides for *how* to give control to applications, and are only applicable after it has been determined *what* level of control is allowable.

In particular, state shared between applications can require fairly high-level semantic guarantees and, hence, any protected interface will be high-level. For example, processes sharing a Unix-flavor file system may require that file modification times be accurate, which involves ensuring disk operations modify file access times. An exokernel is about protection. If such protection is required, providing it in the kernel does not violate exokernel precepts. What the exokernel does say is that non-protection related functionality should be migrated to unprivileged software. In addition, applications not needing these file modification guarantees should not be forced to use them. They should be able to to access the low-level disk in order to implement an alternative file system.

How the kernel is organized. It is easy to read the exokernel principles as an excessive concern for what is in the kernel proper. This is not what is intended. The central question of an exokernel is not about how to organize privileged code. It does not matter whether an exokernel is implemented as a monolithic kernel or as a collection of trusted processes around a micro-kernel. Rather the architecture's central question is how how to make traditionally privileged code unprivileged by limiting the duties of the kernel to just these required for protection.

Acceptable default functionality. An exokernel may provide significant default functionality as long as it can be overridden without compromising either protection or efficiency.

2.4 Visible Resource Revocation

Once resources have been bound to applications, there must be a way to reclaim them and break their bindings. Revocation can either be *visible* or *invisible* to applications. Traditionally, operating systems have performed revocation invisibly, deallocating resources without application involvement. For example, with the exception of some external

pagers [2, 77], most operating systems deallocate (and allocate) physical memory without informing applications. This form of revocation has lower latency and is simpler than visible revocation since it requires no application involvement. Its disadvantages are that library operating systems cannot guide deallocation and have no knowledge that resources are scarce.

An exokernel uses visible revocation for most resources. Even the processor is explicitly revoked at the end of a time slice; a library operating system might react by saving only the required processor state. For example, a library operating system could avoid saving the floating point state or other registers that are not live. However, since visible revocation requires interaction with a library operating system, invisible revocation can perform better when revocations occur very frequently. Processor addressing-context identifiers are a stateless resource that may be revoked very frequently and are best handled by invisible revocation.

Revocation and Physical Naming

The use of physical resource names requires that an exokernel reveal each revocation to the relevant library operating system so that it can relocate its physical names. For instance, a library operating system that relinquishes physical page "5" should update any of its page-table entries that refer to this page. This is easy for a library operating system to do when it deallocates a resource in reaction to an exokernel revocation request. An abort protocol (discussed below) allows relocation to be performed when an exokernel forcibly reclaims a resource.

We view the revocation process as a dialogue between an exokernel and a library operating system. Library operating systems should organize resource lists so that resources can be deallocated quickly. For example, a library operating system could have a simple vector of physical pages that it owns: when the kernel indicates that some page should be deallocated, the library operating system selects one of its pages, writes it to disk, and frees it.

The Abort Protocol

An exokernel must also be able to take resources from library operating systems that fail to respond satisfactorily to revocation requests. An exokernel can define a second stage of the revocation protocol in which the revocation request ("please return a memory page") becomes an imperative ("return a page within 50 microseconds"). However, if a library operating system fails to respond quickly, the bindings need to be broken "by force." The actions taken when a library operating system is recalcitrant are defined by the *abort protocol*.

One possible abort protocol is to simply kill any library operating system and its associated application that fails to respond quickly to revocation requests. We rejected this method because we believe that most programmers have great difficulty reasoning about hard real-time bounds. Instead, if a library operating system fails to comply with the revocation protocol, an exokernel simply breaks all existing bindings to the resource and informs the library operating system.

To record the forced loss of a resource, a *repossession vector* can be used. When an exokernel takes a resource from a library operating system, this fact is registered in the vector and the library operating system receives a "repossession" exception so that it can update any mappings that use the resource. For resources with state, an exokernel can write the state into another memory or disk resource. In preparation, the library operating system can pre-load the repossession vector with a list of resources that can be used for this purpose. For example, it could provide names and capabilities for disk blocks that should be used as backing store for physical memory pages.

Another complication is that an exokernel should not arbitrarily choose the resource to repossess. A library operating system may use some physical memory to store vital bootstrap information such as exception handlers and page tables. The simplest way to deal with this is to guarantee each library operating system a small number of resources that will not be repossessed (*e.g.*, five to ten physical memory pages). If even those resources must be repossessed, some emergency exception that tells a library operating system to submit itself to a "swap server" is required.

2.5 Secure Bindings

One of the primary tasks of an exokernel is to multiplex resources *securely*, providing protection for mutually distrustful applications. To implement protection an exokernel must guard each resource. To perform this task efficiently, an exokernel allows library operating systems to bind to resources using *secure bindings*.

A secure binding is a protection mechanism that decouples authorization from the actual use of a resource. Secure bindings improve performance in two ways. First, the protection checks involved in enforcing a secure binding are expressed in terms of simple operations that the kernel (or hardware) can implement quickly. Second, a secure binding performs authorization only at bind time, which allows management to be decoupled from protection. Application-level software is responsible for many resources with complex semantics (*e.g.*, network connections). By isolating the need to understand these semantics to *bind time*, the kernel can efficiently implement access checks at *access time* because it need no longer involve an application. Simply put, a secure binding allows the kernel to protect resources without understanding them.

Operationally, the one requirement needed to support secure bindings is a set of primitives that application-level software can use to express protection checks. The primitives can be implemented either in hardware or software. A simple hardware secure binding is a TLB entry: when a TLB fault occurs the complex mapping of virtual to physical addresses in a library operating system's page table is performed and then loaded into the kernel (bind time) and then used multiple times (access time). Another example is the packet filter [65], which allows predicates to be downloaded into the kernel (bind time) and then run on every incoming packet to determine which application the packet is for (access time). Without a packet filter, the kernel would need to query every application or network server on every packet reception to determine who the packet was for. By separating protection (determining who the packet is for) from authorization and management (setting up connections, sessions, managing retransmissions, *etc.*) very fast network multiplexing is possible while still supporting complete application-level flexibility.

We use three basic techniques to implement secure bindings: hardware mechanisms, software caching, and downloading application code.

Appropriate hardware support allows secure bindings to be couched as low-level protection operations such that later operations can be efficiently checked without recourse to high-level authorization information. For example, a file server can buffer data in memory pages and grant access to authorized applications by providing them with capabilities for the physical pages. An exokernel would enforce capability checking without needing any information about the file system's authorization mechanisms. As another example, some Silicon Graphics frame buffer hardware associates an ownership tag with each pixel. This mechanism can be used by the window manager to set up a binding between a library operating system and a portion of the frame buffer. The application can access the frame buffer hardware directly, because the hardware checks the ownership tag when I/O takes place.

Secure bindings can be cached in an exokernel. For instance, an exokernel can use a large software TLB [7, 46] to cache address translations that do not fit in the hardware TLB. The software TLB can be viewed as a cache of frequently-used secure bindings.

Secure bindings can be implemented by downloading code into the kernel. This code is invoked on every resource access or event to determine ownership and the actions that the kernel should perform. Downloading code into the kernel allows an application thread of control to be immediately executed on kernel events. The advantages of downloading code are that potentially expensive crossings can be avoided and that this code can run without requiring the application itself to be scheduled. Type-safe languages [9, 75], interpretation, and sandboxing [89] can be used to execute untrusted application code safely [26].

2.6 Methodology Discussion

> You can't learn too soon that the most useful thing about a principle is that it can always be sacrificed to expediency. — W. Somerset Maugham (1874-1965)

Stylistically, exokernel design consists of two different activities: giving applications control of resources, and implementing protection – i.e., making sure they control only their resources. Ideally, a library operating systems has safe, efficient access to anything a privileged OS does.

Exokernel design profits from a deceptively simple shift in goals. Rather than focus on defining the right abstraction, or how to implement it — characteristic of almost all other software design (operating system related or otherwise) — an exokernel architecture concentrates instead deferring these decisions to untrusted software, and then verifying that they are implemented correctly. This requires constructing interfaces that do checking of an operation rather than imperatively deciding how to do it. Thus, algorithmic design becomes a partitioning process of dividing problems into two pieces. The first contains the most “interesting” aspects, which the exokernel leaves to applications. The second part contains the residue that an exokernel must perform to verify correctness. For this partitioning to be practical, checking must be comparatively inexpensive.

As an example, consider the problem of writing cached disk blocks to stable storage in a way that guarantees consistency across reboots. Rather than an exokernel deciding on a particular write ordering and having to struggle with the associated tradeoffs in scheduling heuristics and caching decisions required, it can instead allow the application to construct schedules, retaining for the much simplified task of merely checking that any application schedule gives appropriate consistency guarantees. Application of this methodology enables an exokernel to leave library operating systems to decide on tradeoffs themselves rather than forcing a particular set, a crucial shift of labor.

While this style of design may seem obvious, in practice it has been depressingly easy to slide into "the old way" of deciding how to implement a particular feature rather than asking the unusual question of how one can get out of doing so, safely. A useful heuristic to catch such slips is to note when one is making many tradeoffs. A plethora of tradeoffs almost invariably signals a decision that could be implemented or optimized in different important ways and, thus, should be left to applications.

Because libOSes understand the higher-level semantics of their abstractions, they "just know" many things that an exokernel does not. For example, that related files will likely be clustered together and if one block is fetched, the succeeding eight blocks should be as well. Thus, a variant of the above methodology is designing interfaces where actions do not require justification. As an example, consider the act of fetching a disk block in core. If a file system must present credentials before the kernel will allow the fetch to be initiated, then it faces serious problems doing the prefetching it needs, since it would first have to read in all file directory entries, show each file's capability to the kernel, and only then initiate the fetches. In contrast, by only performing access control when the actual data in the block or read or written, rather then when the block is fetched, these disk reads can be initiated all at once.

Chapter 3

Practice: Applying exokernel principles

> When a man says he approves of something in principle, it means he hasn't the slightest intention of putting it into practice. — Otto von Bismarck (1815-1898)

To illustrate the exokernel principles, this chapter discusses how to export low-level primitives such as exceptions and inter-process communication, and multiplex physical resources such as memory, CPU, and the network. To make this discussion concrete, we draw on examples from two exokernel systems: Aegis [25], and Xok [48]. Aegis is the first exokernel we built. It runs on the MIPS [50] DECstation family. Xok is the second. It runs on the x86 architecture and is the more mature of the two. Construction of these systems spanned four years: two for Aegis, two (and counting) for Xok.

Xok and Aegis differ in important ways, helping to show how the application of exokernel principles changes in the face of different constraints. Most of these differences result from the fact that the x86 and MIPS hardware have radically different architectural interfaces (e.g., software page tables for the MIPS, hardware for x86) and implementation performance (an order of magnitude in favor of the x86).

The implementations we discuss are merely examples of how resources *can be* multiplexed, not how they must be. Other implementations are possible.

Stylistically, the discussion of each resource focuses on two questions: (1) how to give applications control over the resource and (2) how to protect it. Some of the resources we discuss have little to do with protection due to the lack of "sharing." For example, besides memory protection issues, exceptions are contained within a single process. Lack of multiplexing makes protection a non-issue. Others, such as networking, focus more heavily on it.

We provide a quick overview of Aegis and Xok below. The remainder of the chapter discusses how they multiplex resources.

3.0.1 Xok overview

Xok multiplexes most standard machine resources: network, CPU, physical memory, and disk. It currently lacks support for display devices. Its default library operating system, ExOS, provides "Unix in a library." ExOS does not handle some Unix corner cases, but it is not a toy. Many sophisticated applications, such as csh, perl, gcc, and telnet, run without modification. The two main caveats of ExOS are lack of virtual memory paging support (the file buffer cache is paged), and that some shared data structures (such as the global file descriptor table) reside in shared memory and could be corrupted by a malicious process. Neither is an inherent limitation, but simply the result of lack of time. For example, Aegis/ExOS had paging and we are adding more protection to data structures as time allows. Neither does either improve our performance. If anything, lack of paging hurts our experiments since it means that memory that could be used for the file buffer cache is wasted on backing virtual memory. As discussed in Chapter 5 our experiments compensate for this lack of protection.

3.0.2 Aegis overview

Aegis, while older then Xok, is more primitive, lacking solid support for disk and having a more crude approach to access control (flat access control lists). Its version of ExOS (from which that of Xok descends) also is not as developed. Most of our Aegis measurements use micro-benchmarks rather application timings.

The micro-benchmarks discussed in this section compare Aegis and ExOS with the performance of Ultrix4.2 (a mature monolithic UNIX operating system) on the same hardware. While Aegis and ExOS do not offer the same level of functionality as Ultrix, we do not expect these additions to cause large increases in our timing measurements.

Ultrix, despite its poor performance relative to Aegis, is *not* a poorly tuned system; it is a mature monolithic system that performs quite well in comparison to other operating systems [69]. For example, it performs two to three times better than Mach 3.0 in a set of I/O benchmarks [67]. Also, its virtual memory performance is approximately twice that of Mach 2.5 and three times that of Mach 3.0 [5].

Our measurements were taken on a DECstation5000/125 (25MHz), with an R3000 processor and a SPECint92 rating of 25.

3.1 Multiplexing Physical Memory

Physical memory is one of the simplest resources to multiplex. When a library operating system allocates a physical memory page, the exokernel records the owner and permissions (e.g., read and write) of the allocating process. The owner of a page has the power to change the capabilities associated with it, to share it, and to deallocate it.

To ensure protection, the exokernel guards every access to a physical memory page by requiring that the capability be presented by the library operating system requesting access. If the capability is insufficient, the request is denied. Typically, the processor contains a TLB, and the exokernel must check memory capabilities when a library operating system attempts to enter a new virtual-to-physical mapping. To improve library operating system performance by reducing the number of times secure bindings must be established, an exokernel may cache virtual-to-physical mappings in a large software TLB.

If the underlying hardware defines a page-table interface, then an exokernel must guard the page table instead of the TLB. Although the details of how to implement secure memory bindings will vary depending on the details of the address translation hardware, the basic principle is straightforward: privileged machine operations such as TLB loads and DMA must be guarded by an exokernel. As dictated by the exokernel principle of exposing kernel book-keeping structures, the page table should be visible (read only) at application level.

To reclaim a page, an exokernel must change the associated capabilities, mark the resource as free, and remove all bindings. In the example of physical memory, bindings include TLB mappings and any queued DMA requests.

3.1.1 Aegis: application virtual memory

The MIPS architecture has software defined page tables. Thus, in accordance with deferring management to applications, Aegis allows applications to define their own page tables. We look at two issues in supporting application-level virtual memory: bootstrapping and efficiency.

To bootstrap the virtual naming system, there must be support for translation exceptions on both application page-tables and exception code. Aegis provides a simple bootstrapping mechanism through the use of a small number of guaranteed mappings. A miss on a guaranteed mapping will be handled automatically by Aegis. This organization frees the application from dealing with the intricacies of boot-strapping TLB miss and exception handlers, which can take TLB misses. To implement guaranteed mappings efficiently, an application's virtual address space is partitioned into two segments. The first segment holds normal application data and code. Virtual addresses in this segment can be "pinned" using guaranteed mappings. Typically libOSes pin exception handling code and page-tables.

On a TLB miss, the following actions occur:

1. Aegis checks which segment the virtual address resides in. If it is in the standard user segment, the exception is dispatched directly to the application. If it is in the second region, Aegis first checks to see if it is a guaranteed mapping. If so, Aegis installs the TLB entry and continues; otherwise, Aegis forwards it to the application.

2. The application looks up the virtual address in its page-table structure and, if the access is not allowed raises the appropriate exception (*e.g.*, "segmentation fault"). If the mapping is valid, the application constructs the appropriate TLB entry and its associated capability and invokes the appropriate Aegis system routine.

3. Aegis checks that the given capability corresponds to the access rights requested by the application. If it does, the mapping is installed in the TLB and control is returned to the application. Otherwise an error is returned.

4. The application performs cleanup and resumes execution.

In order to support application-level virtual memory efficiently, TLB refills must be fast. To this end, Aegis caches TLB entries (a form of secure bindings) in the kernel by overlaying the hardware TLB with a large software TLB (STLB) to absorb capacity misses [7, 46]. On a TLB miss, Aegis first checks to see whether the required mapping is in the STLB. If so, Aegis installs it and resumes execution; otherwise, the miss is forwarded to the application.

The STLB contains 4096 entries of 8 bytes each. It is direct-mapped and resides in unmapped physical memory. An STLB "hit" takes 18 instructions (approximately one to two microseconds). In contrast, performing an upcall to application level on a TLB miss, followed by a system call to install a new mapping is at least three to six microseconds more expensive.

As dictated by the exokernel principle of exposing kernel book-keeping structures, the STLB can be mapped using a well-known capability, which allows applications to efficiently probe for entries.

Using the control given by libOS-defined page tables ExOS has a variety of different page table structures [1], high-performance network paging, and application-specific page coloring (for improved cache performance).

No other protected operating system architecture allows this degree of freedom.

3.1.2 Xok: hardware-defined page tables

Unlike the MIPS architecture, the x86 architecture on which Xok runs defines the page-table structure. Since x86 TLB refills are handled in hardware, this structure cannot be overridden by applications. However, applications are able to control all hardware-defined per-page attributes, including protection levels (read-only, writable, execute-only) and caching. Additionally, each valid page-table entry contains three software-defined bits, which Xok gives over to application control. Invalid entries (those without the "present bit" set) can contain any value the library operating system wishes. Example uses of this location are to store the names of disk blocks used as backing store or even network addresses for remote pages.

Since the hardware does not verify that the physical page of a translation can be mapped by a process, applications are prevented from directly modifying the page table and must instead use system calls. Although these restrictions make Xok less extensible than Aegis, they simplify the implementation of libOSes (see Chapter 7 for more discussion).

Like Aegis, Xok allows efficient and powerful virtual memory abstractions to be built at the application level. It does so by exposing the capabilities of the hardware (e.g., all MMU protection and dirty bits) and exposing many kernel data structures (e.g., free lists, inverse page mappings). Xok's low-level interface means that paging is handled by applications. As such, it can be done from disk, across the network, or by data regeneration. Additionally, applications can readily perform per-page transformations such as compression, verification of contents using digital signatures (to allow untrusted nodes in a network to cache pages), or encryption.

3.2 Multiplexing the Network

This subsection discusses how to give applications control over the network, and how to implement protection. There are two primary requirements for efficient networking. First, elimination of data copies, which comes from both giving applications access to any scatter-gather DMA provided by the hardware and allowing them to direct where messages are placed. Second, tight coupling to interrupt events, primarily, message reception and timer interrupts. In order to initiate low-latency responses to messages, application messaging code must be able to run quickly after message arrival. Aegis performs this by downloading application code into the kernel and running it in the interrupt handler [92]. Xok, due to the vast relative increase in the ratio of processor speed to network latency (about a factor of five to ten for small messages), does not need to eliminate boundary crossings, and simply yields to the receiving application. Timer interrupts are needed to build efficient retransmission timers, which are needed to implement fast reliable messaging on unreliable network hardware.

To enable application resource management, an exokernel dislocates operating system code into libraries. However, there are important differences between a library's execution context and that of the kernel. One of the challenges in an exokernel is recapturing these aspects. Tight coupling of libOS code to events can be viewed as an example of this. Compared to Ultrix, Aegis's provision of efficient access to these two abilities gives library operating systems a factor of ten performance improvement in round trip Ethernet times [25].

[1] Illustrative of the customizability of an exokernel system, Tom Pinckney while still an undergraduate was able to implement a new page table structure in a week, while taking his final exams. As a testament to the difficulty of modifying current operating systems, the proposers of this page table structure were only able to simulate it [83].

Protection for networking is answering the question: given a message, who owns it? On a connection-oriented network, answering this question is easy: whichever connection (or flow) it belongs to. Binding of application to flow can happen at connection initiation, removing the need to make decisions based on packet contents. An example of a hardware-based mechanism is the use of the virtual circuit in ATM cells to securely bind streams to applications [23]. However, answering this question is more difficult on a connectionless network such as Ethernet, where message ownership requires understanding header semantics. Since the exokernel has dislocated all network protocol code that understands packet semantics into library operating systems it lacks the information necessary to decide which application owns what message.

To solve this problem, an exokernel requires that networking libraries download packet filters [65] to select the messages they want. [2] Conceptually, every filter is invoked on every arriving packet and returns "accept" (the filter wants the message) or "reject" (the filter does not want the message). The operating system thus need not understand the actual bits in a message in order to bind an arriving packet to its owner.

For protection, the exokernel must ensure that that a filter does not "lie" and accept packets destined to another process. To prevent this theft we intentionally designed our filter language to make "overlap" detection simple. (Alternatively, simple security precautions such as only allowing a trusted server to install filters could be used to address this problem.) Finally, we ensure fault isolation through a combination of language design (to bound runtime) and runtime checks (to protect against wild memory references and unsafe operations).

Both Aegis and Xok use packet filters, because our current network does not provide hardware mechanisms for message demultiplexing. One challenge with a language-based approach is to make filters fast. Traditionally, packet filters have been interpreted, making them less efficient than in-kernel demultiplexing routines. One of the distinguishing features of the packet filter engine used by our prototype exokernel is that it compiles packet filters to machine code at runtime, increasing demultiplexing performance by more than an order of magnitude [28, 31]. [3]

Sharing the network interface for outgoing messages is easier. Messages are simply copied from application space into a transmit buffer. In fact, with appropriate hardware support, transmission buffers can be mapped into application space just as easily as physical memory pages [23].

An exokernel defers message construction to applications. A protection problem this creates is how to prevent applications from "spoofing" other applications by sending bogus messages. Our two exokernel systems do not prevent this attack, since they were designed for an insecure networking environment. However, within the context of a trusted network, an exokernel could use "inverse" packet filters to reject messages that do not fit a specified pattern (or, alternatively, reject messages that do). The techniques that worked for DPF could be applied here as well.

3.3 Multiplexing the CPU

Control over the CPU involves the ability to: (1) allocate and share CPU time-slices (similar to other resources), (2) yield a time-slice to a specific named process and (3) receive notification (e.g., via an upcall) when time-slices begin and end. The attributes of time slices are quantity and "timeliness." Applications should be able to allocate specific time slices to control either of these attributes. Finally, the notion of "process" should be low-level, consisting of the information required by hardware, such as program counters to vector machine exceptions to, and protection-required information such as a binding between the process and a principal.

The following discussion makes these rules concrete by considering Aegis' treatment of the CPU and its process and exception facilities.

3.3.1 Aegis and Xok CPU multiplexing

Both Xok and Aegis multiplex the CPU in the same way. They represent the CPU as a linear vector, where each element corresponds to a time slice. Time slices are partitioned at the clock granularity and can be allocated in a manner similar to physical memory. Scheduling is done "round robin" by cycling through the vector of time slices. A crucial property of this representation is *position*, which encodes an ordering and an approximate upper bound on when

[2] Packet filters originated to solve the equivalent difficulty brought about by another flavor of operating system code motion — microkernels — which, because they moved networking code from out of the operating system into privileged servers, rendered an operating system too ignorant to demultiplex packets.

[3] However, recently we have experimented with the use of an aggressive interpreter that, by preprocessing filters and exploiting "super-operator" instructions, runs roughly within a factor of four of hand-tuned code — a perfectly adequate speed given that demultiplexing is coupled to a high-latency I/O event (packet reception).

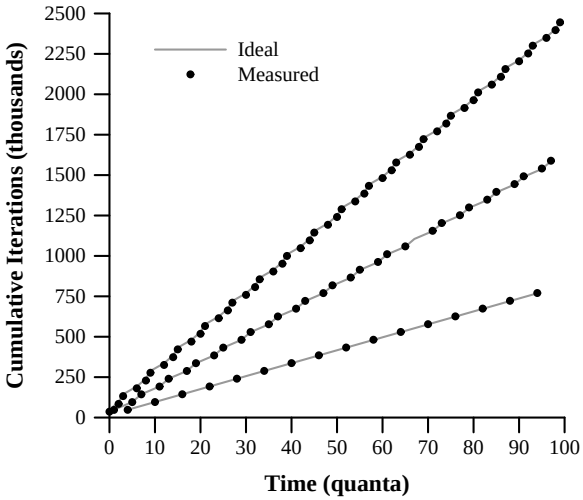

Figure 3-1: Application-level stride scheduler.

the time slice will be run. Position can be used to meet deadlines and to trade off latency for throughput. For example, a long-running scientific application could allocate contiguous time slices in order to minimize the overhead of context switching, while an interactive application could allocate several equidistant time slices to maximize responsiveness.

The kernel provides a yield primitive to donate the remainder of a process' current time slice to another (specific) process. Applications can use this simple mechanism to implement their own scheduling algorithms.

Timer interrupts denote the beginning and end of time slices, and are delivered in a manner similar to exceptions (discussed below): a register is saved in the "interrupt save area," the exception program counter is loaded, and the kernel jumps to user-specified interrupt handling code with interrupts re-enabled. The application's handlers are responsible for general-purpose context switching: saving and restoring live registers, releasing locks, *etc*. This framework gives applications a large degree of control over context switching. For example, because it notifies applications of clock interrupts, and gives them control over context switching's state saving and restoration, it can be used to implement scheduler activations [4].

Fairness is achieved by bounding the time an application takes to save its context: each subsequent timer interrupt (which demarcates a time slice) is recorded in an excess time counter. Applications pay for each excess time slice consumed by forfeiting a subsequent time slice. If the excess time counter exceeds a predetermined threshold, the environment is destroyed. In a more friendly implementation, the kernel could perform a complete context switch for the application.

This simple scheduler can support a wide range of higher-level scheduling policies. As we demonstrate below, an application can enforce proportional sharing on a collection of sub-processes.

Extensible Schedulers on Aegis

Using the Aegis yield primitive and control over time slices, we have built an application-level scheduler that implements *stride scheduling* [91], a deterministic, proportional-share scheduling mechanism that improves on recent work [90]. The ExOS implementation maintains a list of processes for which it is responsible, along with the proportional share they are to receive of its time slice(s). On every time slice wakeup, the scheduler calculates which process is to be scheduled and yields to it directly.

We measure the effectiveness of this scheduler by creating three processes that increment counters in shared memory. The processes are assigned a 3:2:1 relative allocation of the scheduler's time slice quanta. By plotting the cumulative values of the shared counters, we can determine how closely this scheduling allocation is realized. As can be seen in Figure 3-1, the achieved ratios are very close to idealized ones, showing that applications can effectively control scheduling.

It is important to note that there is nothing special about this scheduler either in terms of privileges (*any* application can perform identical actions) or in its complexity (the entire implementation is less than 100 lines of code). As a result, any application can easily manage processes.

3.3.2 Aegis Processor Environments

An Aegis processor environment is a structure that stores the information needed to deliver events to applications. All resource consumption is associated with an environment because Aegis must deliver events associated with a resource (such as revocation exceptions) to its designated owner.

Four kinds of events are delivered by Aegis: exceptions, interrupts, protected control transfers, and address translations. Processor environments contain the four contexts required to support these events:

Exception context: for each exception, an exception context contains a program counter for where to jump to and a pointer to physical memory for saving registers.

Interrupt context: for each interrupt an interrupt context includes a program counters and register-save region. In the case of timer interrupts, the interrupt context specifies separate program counters for start-time-slice and end-time-slice cases, as well as status register values that control co-processor and interrupt-enable flags.

Protected Entry context: a protected entry context specifies program counters for synchronous and asynchronous protected control transfers from other applications. Aegis allows any processor environment to transfer control into any other; access control is managed by the application itself.

Addressing context: an addressing context consists of a set of *guaranteed mappings*. A TLB miss on a virtual address that is mapped by a guaranteed mapping is handled by Aegis. Library operating systems rely on guaranteed mappings for bootstrapping page-tables, exception handling code, and exception stacks. The addressing context also includes an address space identifier, a status register, and a tag used to hash into the Aegis software TLB. To switch from one environment to another, Aegis must install these three values.

These are the event-handling contexts required to define a process. Each context depends on the others for validity: for example, an addressing context does not make sense without an exception context, since it does not define any action to take when an exception or interrupt occurs.

3.3.3 Implementing Unix Processes on Xok

The *process map* maps UNIX process identifiers to Xok environment numbers using a shared table. The *process table* records the process identifiers of each process, that of its parent, the arguments with which the process was called, its run status, and the identity of its children. The table is partitioned across application-reserved memory of Xok's environment structure, which is mapped readable for all processes and writeable for only the environment's owning process. ExOS uses Xok's IPC to safely update parent and child process state. The UNIX *ps* (process status) program is implemented by reading all the entries of the process table.

UNIX provides the *fork* system call to duplicate the current process and *exec* to overlay it with another. *Exec* is implemented by creating a new address space for the new process, loading on demand the disk image of the process into the new address space, and then discarding the address space that called *exec*. Implementing fork in a library is peculiar since it requires that a process create a replica of its address space and state *while it is executing*. To make fork efficient, ExOS uses copy-on-write to lazily create separate copies of the parent's address space. ExOS scans through its page tables, which are exposed by Xok, marking all pages as copy-on-write except those data segment and stack pages that the *fork* call itself is using. These pages must be duplicated so as not to generate copy-on-write faults while running the *fork* and page fault handling code. Groups of page table entries are updated at once by batching system calls to amortize the system call overhead over many updates.

3.4 Exposing Machine Events

An exokernel gives applications low-level efficient access to machine events such as exceptions and interrupts (which we have already seen in the context of networking and the CPU). It also provides primitives for inter-domain calls, which safely change the program counter in one process to an agreed upon value in another's address space. Below, we discuss Aegis' mechanisms for exceptions and control transfer, along with Xok's facilities for interprocess communication.

3.4.1 Aegis Exceptions

Aegis dispatches all hardware exceptions to applications (save for system calls) using techniques independently developed but similar to those described in Thekkath and Levy [87]. To dispatch an exception, Aegis performs the following actions:

1. It saves three scratch registers into an agreed-upon "save area." (To avoid TLB exceptions, Aegis does this operation using physical addresses.)

2. It loads the exception program counter, the last virtual address that failed to have a valid translation, and the cause of the exception.

3. It uses the cause of the exception to perform an indirect jump to an application-specified program counter value, where execution resumes with the appropriate permissions set (*i.e.*, in user-mode with interrupts re-enabled).

After processing an exception, applications can immediately resume execution without entering the kernel. Ensuring that applications can return from their own exceptions (without kernel intervention) requires that all exception state be available for user reconstruction. This means that all registers that are saved must be in user-accessible memory locations.

Currently, Aegis dispatches exceptions in 18 instructions (1.5 microseconds on our 25MHz DECstation5000/125), roughly a factor of a hundred faster than Ultrix, a monolithic OS running on the same hardware, and over five times faster than the most highly-tuned implementation in the literature. Mainly this improvement comes from the two facts that (1) exokernel primitives are low-level, which means applications do not pay for unnecessary functionality, and (2) the exokernel concentrates on doing a few things well. For example, Aegis is small, and does not need to page its data structures. Because, all kernel addresses are physical they never miss in the TLB, and Aegis does not have to separate kernel TLB misses from the more general class of exceptions in its exception demultiplexing routine.

Fast exceptions enable a number of intriguing applications: efficient page-protection traps can be used by applications such as distributed shared memory systems, persistent object stores, and garbage collectors [5, 87]. Exception times are discussed further in [25].

3.4.2 Aegis Protected Control Transfers

Aegis provides a *protected control transfer* mechanism as a substrate for efficient implementations of inter-process compunction (IPC) abstractions. This mechanism illustrates the principle of not encapsulating primitives within high-level abstractions. It consists of the minimum operations needed to transfer execution from one process to another in a safe way: it changes the program counter to an agreed-upon value in the callee, donates the current time slice to the callee's processor environment, and installs the required elements of the callee's processor context (addressing-context identifier, address-space tag, and processor status word).

Aegis provides two forms of protected control transfers: *synchronous* and *asynchronous*. The difference between the two is what happens to the processor time slice. Asynchronous calls donate only the remainder of the current time slice to the callee. Synchronous calls donate the current time and all future instantiations of it; the callee can return the time slice via a synchronous control transfer call back to the original caller. Both forms of control transfer guarantee two important properties. First, to applications, a protected control transfer is atomic: once initiated it will reach the callee. Second, Aegis will not overwrite any application-visible register. These two properties allow the large register sets of modern processors to be used as a temporary message buffer [17].

Currently, our synchronous protected control transfer operation takes 30 instructions (1.4 microseconds on a 25MHz DECstation5000/125) [25]. Roughly ten of these instructions are used to distinguish the system call "exception" from other hardware exceptions on the MIPS architecture. Setting the status, co-processor, and address-tag registers consumes the remaining 20 instructions, and could benefit from additional optimizations. Because Aegis implements the minimum functionality required for any control transfer mechanism, applications can efficiently construct their own IPC abstractions. As an example, we have implemented a "trusted" local remote procedure call for use between a client and a server that the client trusts to save and restore callee-saved registers. By eliminating the need to save all registers, trusted LRPC improves performance by roughly a factor of two.

Hardware differences make comparing our numbers to others in the literature difficult. We attempt an extremely crude comparison of our protected control transfer operation to the equivalent operation on L3 [57], the fastest published result by normalizing the IPC on both systems using SPECint92 rating of Aegis's DEC5000 and L3's 486 (16.1 vs. 30.1). Aegis's trusted control transfer mechanism is 6.6 times faster than the scaled time for L3's RPC mechanism.

Given the difference in architectures, this scaling should be treated as a "back of the envelop" computation. The main conclusion to draw is that Aegis' control transfer is fast.

3.4.3 Implementing Unix IPC on Xok

UNIX defines a variety of interprocess communication primitives: signals (software interrupts that can be sent between processes or to a process itself), pipes (producer-consumer untyped message queues), and sockets (differing from pipes in that they can be established between non-related processes, potentially executing on different machines).

Signals are layered on top of Xok IPC. Pipes are implemented using Xok's *software regions*, which provide sub-page memory protection, coupled with a "directed yield" to the other party when it is required to do work (i.e., if the queue is full or empty). Sockets communicating on the same machine are currently implemented using a shared buffer.

Inter-machine sockets are implemented through user-level network libraries for UDP and TCP. The network libraries are implemented using Xok's timers, upcalls, and packet rings, which allow protected buffering of received network packet,

3.5 Discussion

This chapter has discussed how an exokernel can safely export a variety of resources — physical memory, the CPU, the network, and hardware events — to applications. The low-level of the exokernel interface allows library operating systems, working above the exokernel interface, to implement higher-level abstractions and define special-purpose implementations that best meet the performance and functionality goals of applications. This organization allows the redefinition of fundamental operating system abstractions by simply changing application-level libraries. Furthermore, these different versions can co-exist on the same machine and are fully protected by Aegis.

The next chapter presents the resource we have found most challenging, disk, along with our solution to it (and five failed solutions).

Chapter 4

The Hardest Multiplexing Problem: Disk

An exokernel must provide a means to safely multiplex disks among multiple library file systems (libFSes). Each libOS contains one or more libFSes. Multiple libFSes can be used to share the same files with different semantics. In addition to accessing existing files, libFSes can define new on-disk file types with arbitrary meta data formats. An exokernel must give libFSes as much control over file management as possible while still protecting files from unauthorized access. It therefore cannot rely on simple-minded solutions like partitioning to multiplex a disk: each file would require its own partition.

To allow libFSes to perform their own file management, an exokernel stable storage system must satisfy four requirements. First, creating new file formats should be simple and lightweight. It should not require any special privilege. Second, the protection substrate should allow multiple libFSes to safely share files at the raw disk block and meta data level. Third, the storage system must be efficient—as close to raw hardware performance as possible. Fourth, the storage system should facilitate cache sharing among libFSes, and allow them to easily address problems of cache coherence, security, and concurrency.

The goal of disk multiplexing is to make untrusted library file systems (libFSes) as powerful as privileged file systems. As listed above, there are a number of engineering challenges to reaching this goal. The hardest challenge by far, however, is access control: i.e., answering the deceptively simple question “who can use a given disk block?” Inventing a satisfactory solution to this problem took us three years and four systems. This chapter describes how we multiplex stable storage, both to show how we address these problems and to provide a concrete example of the exokernel principles in practice. We first describe out solution to efficient access control. An interesting aspect of this solution is that it uses the libFSes own data structures to track what the libFS owns. Another is the the code verification technique we invented to do such re-use without impinging on libFS flexibility. We then give an overview XN, Xok's extensible, low-level in-kernel stable storage system. We also describe the general interface between XN and libFSes and present one particular libFS, C-FFS, the co-locating fast file system [37].

4.1 Efficient, fine-grained disk multiplexing

To provide protection, the operating system must force each libFS to only access those disk blocks for which they have access rights. For flexibility and speed an exokernel provides fine-grained protection: i.e., at the level of disk blocks rather than partitions. To do so we must answer the simple question “who can use this block?” Doing so requires constructing a mapping of each disk block to every principal that can use it, which turns out to be quite difficult [48] mainly because it requires building the moral equivalent of a file system. For example, the mapping table will be enormous (since its size is proportional to the size of disk) and, hence, not fit in main memory. As a result, managing it as it is is moved between both memory and disk requires solving all the standard file systems issues — deciding which pieces to cache, prefetch and evict, how to perform allocation of the table on disk, how to reconstruct the table's state after a system crash, etc. Obviously, if an exokernel already implements a file system, libFSes built on top of it will have little freedom. Fortunately, there are a series of insights that can allow us to trust the libFS itself to efficiently track the blocks it owns, thereby eliminating the need for duplicate bookkeeping.

The first insight is to notice that a correct libFS will already track what disk blocks it owns; doing so efficiently is, after all, one of the main purposes of a file system. Therefore, any bookkeeping an exokernel does is redundant. Thus,

if it can reuse the bookkeeping data structures of the libFS then we can eliminate this redundancy. [1] For example, a Unix file system tracks what blocks are associated with it (and what principals are allowed to use them) using "meta data" consisting of directory blocks, inodes, as well as single, double, and triple indirect blocks. In other words, everything the kernel needs to track to perform access control. If it can reuse libFS meta data, then, it does not need to perform duplicate tracking of what blocks are associated with which libFS.

Our new problem is that reusing libFS bookkeeping structures requires that we understand them, both so that we can force them to be correct, and so that we can extract ownership information from them. One traditional solution would be to provide a fixed set of components (e.g., pointers to disk blocks, disk block extents, etc.) from which libFSes could build their meta data from. However, creating a set of universal building blocks for file system meta data is so hard as to be infeasible: file system research is still an active area even after three decades. A component set capable of describing all results of this research does not seem possible. Instead, we solve this problem by having the libFS interpret its meta data for us in a way that we can test for correctness.

Determinism + induction = verification

Our struggle has two conflicting pieces: flexibility (in that we want to allow client meta data to have any "syntax" and semantics whatsoever, including those that we have not anticipated) and validity (in that we require that they do so correctly, even when we don't understand its representation). *Untrusted deterministic functions* (UDFs) let us achieve both goals: clients can use any possible meta data representation, while we still can verify correctness, no matter how mysterious the representation they chose.

Flexibility comes from adding a layer of indirection: rather than have meta data built from components the operating system understands, libFSes provide an "interpreter" function, owns, for each meta data type. Given an instance of that type, the owns UDF produces the set of blocks that that instance controls: owns: (meta) → set of blocks controlled by meta. For example, a Unix file system will provide an owns function for each of the types of meta data it uses — inodes, directory blocks, indirect nodes, double indirect nodes, triple indirect nodes, etc. Thus, when the the operating system needs to know what disk blocks a piece of meta data controls, it simply runs that meta data through its associated owns function. Since owns is written in a Turing complete language (in our implementation a pseudo-assembly language) it can describe any possible file system meta data layout.

At this point, while we have flexibility, guarantees of correctness have become more challenging. owns is written in general purpose code. Verification (i.e., that owns correctly implements its specification) is beyond current formal methods. The UDF techniques discussed in the remainder of the chapter are an online alternative that is both simple and robust.

UDFs start from the fact that if a deterministic Turing machine terminates and produces an output O then, given the same initial state and input, it will always terminate and produce O. The key to verification is that determinism gives persistence: once owns(meta) produces a valid result, it will always do so, until meta is allowed to change. Without persistence past tests can tell us nothing about present or future behavior.

To make UDFs deterministic we must guarantee three conditions:

1. The instructions the UDF executes have deterministic semantics.

2. No information can be allowed to flow into the universe of the UDF and its state. For example, on each invocation of the UDF all registers and stack locations must be initialized to the same values, the UDF cannot use self-modifying code, call non-deterministic functions (such as a "get time of day" routine), have access to cycle counters, etc.

3. All state that persists across a UDF's invocations is visible to the verifier so that it can retest the UDF when this state is modified.

We assume that UDFs are made of immutable code and protected data and that the execution of this code has been made "safe" to protect the testing evaluator from malice: the code is preventing from corrupting the evaluator's state, looping "too long," etc. The implementation discussed in this section guarantees these conditions using a safe interpreter. A trusted compiler could also be used.

[1] Note that this insight applies to all all areas that deal with protection, not just file systems: correct applications track what resources they own. Future work will involve exploiting this fact more aggressively.

Determinism gives persistence and, as a result, allows us to use induction to verify UDF correctness. Inductive testing has two phases: initialization and modification. Initialization tests that the UDF's initial state is a valid one. Modification tests that when a UDF's state is mutated, the mutation leaves the UDF in a valid state.

Thus, to trust libFS meta data and its associated owns function, we need to check that owns(meta) produces a valid output when meta is initialized, and then retest it after each modification. Once we can force all libFS meta data to produce valid results, we have accomplished the goal of this section: libFSes can now track their disk blocks without the kernel having to duplicate this bookkeeping.

More specifically, initialization checks that when the libFS allocates a piece of meta data, it should not control any disk blocks:

```
% verify owns(meta) is in a valid initial state:
% i.e., meta should control no disk blocks.
proc initialize(meta)
  if owns(meta) != {}
    error "Bogus initial state!";
end;
```

If owns(meta) does not yield the empty set then we know that either owns is incorrect or meta is not properly initialized. In either case we reject meta. (We do not actually care if owns itself is correct — i.e., that it works on all possible inputs — just that it work on the current set of used inputs.) Otherwise, we can accept meta. Because owns is deterministic we are guaranteed that until meta is modified owns(meta) will always produce the empty set.

Next, when a file system wants to modify its meta data (say to allocate a disk block to a file) we verify that meta goes from its current valid state to a new valid state, where, because of determinism, it must remain.

Allocation is thus:

```
% Give meta control of disk block b
proc allocate(meta, b)
  old_set = owns(meta)  % record original state
  <let libFS scribble on meta as it will>
  new_set = owns(meta)    % record new state

  % check that owned set grew by exactly b.
  if new_set != old_set U { b }
    error "Bogus modification!";
end;
```

Because owns is deterministic we can find out its current state by simply querying it. This is the crucial part of the process that frees us from duplicating any bookkeeping data structures. Using that current state we can ensure that the modification goes to a valid next state with a straightforward process of computing set union and equality.

Deallocation is similar to allocation:

```
% Deallocate disk block b from meta
proc deallocate(meta, b)
  old_set = owns(meta);      % record original state
  <let libFS scribble on meta as it will>
  new_set = owns(meta);          % record new state
  % check that owned set shrunk by exactly b.
  if set_diff(new_set, b) != old_set
    error "Bogus modification!";
end;
```

While we must run owns in a safe evaluation context after meta has been modified, any subsequent invocation of owns(meta) can run without safety checks since owns is deterministic and, after testing, we know that it executes safely on this value of meta. In the pseudo-code above, while new_set = owns(meta) must be run in a safe context, the statement old_set = owns(meta) can run unchecked. One way to look at this fact is that halting problem is trivial for deterministic Turing machines that one knows have halted in the past.

Naively, UDF induction might appear to be nothing more than a simplistic application of testing pre- and post-conditions. The crucial difference is that the pre-condition is supplied by the untrusted UDF implementor, who has been rendered a trustworthy partner through determinism.

Determinism and XN's trusted set implementation (used to check UDF output) prevents UDFs from "cheat" and producing bogus output after testing.

At this point we can incrementally verify that owns implements its specification correctly. As a result, the operating system can now rely on potentially malicious libFSes to track what disk blocks they own, in ways that the operating system does not understand, while still being able to guarantee correctness. Figure 4-1 sketches how the above verification steps are embedded in the context of a block allocation system call.

Because our approach merely verifies *what* a UDF did rather than *how* it did so it is both more robust and simpler than traditional verification approaches. Only a handful of lines of pseudo-code are required to verify the correctness of code that is written in an untyped, general-purpose assembly language that allows pointers (including casts between integers and pointers), aliasing, stores, dynamic memory allocation, arbitrary loops, and unstructured control flow. Automated theorem provers, in contrast, are both complex and unable to verify such code.

4.1.1 Efficiency: State Partitioning

Since an instance of meta data can control a large number of blocks, enumerating all blocks after each modification can be inefficient. (For example, the Unix "indirect block" in Figure 4-2 controls up to 1024 disk blocks.) Fortunately, the computation to produce a given element in a UDF's owns set typically uses only a limited portion of the UDF's state. This skewed usage can be exploited by *state partitioning*, which at a high level allows UDFs to partition their associated state into sets of disjoint ranges, where each set contains all the state needed to compute a subset of owned blocks (much smaller in size than the own set controlled by the meta data). Our inductive steps proceed as before, with the two changes that at initialization we need to verify that the partitions are non-overlapping and at modification that the libFS only modifies state in the indicated partition.

More concretely, we number partitions from 0 to n-1 and, on initialization, verify that these partitions do not overlap (calling the access function for a partition id gives the set of bytes in that partition):

```
set = {};
foreach partition id
  if(set_overlap(set, access(id))
     error "Partitions cannot overlap!";
  set = set U access(id);
```

Allocation checks that the given partition id is valid, and then adds the block to the set:

```
if id ≥ the number of partitions
  error "Bogus partition!";

old_set = owns(id);
<let libFS modify state in access(id)>
new_set = owns(id);

if old_set U db != new_set
   error "Bogus modification";
```

```
/*
 * C  pseudo code sketch of how to allocate block 'req' and stuff a pointer to
 * it in the parent.   For simplicity we assume meta is BLOCKSIZE big and that
 * setting a block to all zeros is a valid initialization.
 */
int sys_alloc_blk(blk_t req, blk_t parent, void *new_meta) {
  set_t old_owns, new_owns;
  set_t (*owns)(void *);      /* pointer to owns function */
  void *old_meta, *kid_meta;

  if(!isfree(disk_freemap, req))              /* Is requested block is on the freelist? */
    return NOT_FREE;
  if(!can_read(new_meta, BLOCKSIZE))          /* Is [new_meta,new_meta+BLOCKSIZE) valid memory? */
    return NOT_VALID;
  if(!(old_meta = buffer_cache_lookup(parent))) /* Is parent in core? */
    return NOT_IN_CORE;
  if(!can_write(parent))                      /* Can the current process write to parent? */
    return CANNOT_ACCESS;

  owns = owns_lookup(parent);                 /* get owns function for parent. */

  owns_old = owns(old_meta);                  /* compute current and potential owned sets. */
  owns_new = safe_eval(owns(new_meta));

  /* if owns(new_meta) did an illegal operation will be nil. */
  if((owns_new = safe_eval(owns_function(new_meta)))
    return ILLEGAL_OP;

  /* compare: old U req = new */
  if(!set_equal(owns_new, set_union(owns_old, req)))
    return BOGUS_UPDATE;

  /* allocate a buffer cache entry to hold kid. */
  if(!(kid_meta = buffer_cache_alloc(req)))
    return CANNOT_ALLOC;

  memcpy(old_meta, new_meta, BLOCKSIZE);      /* overwrite parent with new value. */
  memset(kid_meta, 0, BLOCKSIZE);             /* initialize kid buffere cache entry to all zeros. */
  set_dirty(parent);                          /* ensure parent will be flushed back to disk. */

  return SUCCESS;
}
```

Figure 4-1: Implementation sketch of a disk block allocation system call that uses UDFs to let untrusted file systems track the blocks they control.

```
% Unix file system indirect block.
struct indirect_block {
    unsigned blocks[1024];
};

% UDF that returns the (singleton) set of
% [offset, nbyte) tuples of all bytes in partition i.
set indirect_access(int id) {
  return { [sizeof(unsigned) * id, sizeof(unsigned)] };
}

% return number of partitions in an indirect block.
int indirect_npartitions(void) {
  return 1024;
}
```

Figure 4-2: Unix indirect block and its associated state partitioning UDFs, indirect_access and indirect_npartitions. The UDFs are "constant" in that they do not use the meta data block to compute partitions. Non-constant UDFs can be used as long as they are retested when the state they depend on has been modified.

Partitioning of state is controlled by UDFs, since it is their implementor that knows the natural partitions of the problem. For instance, in the simple case of the Unix indirect block given in Figure 4-2, which is simply a vector of 1024 disk block pointers, a partition corresponds to the 32 bits in which a single block pointer is stored. In most situations we have encountered state partitioning can reduce output sets to singleton entries.

In some sense, we have already been doing coarse-grain state partitioning, since we only rerun a UDF when the disk block it is associated with is modified rather than requiring that a file system have a single UDF that is run on modification to any block on disk. Partitioning simply takes this subdivision to finer levels.

A natural question is how to track the state in each partition. Our original implementation use a static partition table for each type, where partition_table(i) gave the set of state sets associated with partition i. Such tables can be large. In the case of file systems, this size is mitigated by the fact that file systems have only a few templates (one for each meta data type). Unfortunately, partitioning via a data structure means that, in many cases, the partitioning cannot be adjusted dynamically, as is required for dynamically resized structures. Of course, the obvious solution is to use UDFs: they were developed, after all to describing meta data layout. This solution turns out to be quite workable, and has significantly reduced the size of templates that we use to store file system meta data descriptions. [2]

To simplify writing UDFs we can refine the partition scheme to make the enumeration of a partition's "read set" implicit rather than explicit. Instead of requiring the UDF client supply an access function, we instead have the safe UDF evaluator trace the memory locations examined by a particular UDF invocation. Initialization checks that these read sets do not overlap by tracing owns on each partition id rather than calling access. Modification checks that the read set of the modified meta data on a given id is the same as the original read set. One complication of this model is handling the case where the read set grows or shrinks. Due to space limitations we elide a discussion of how to do this gracefully; interested readers can refer to [30].

4.1.2 More sophisticated partitioning

Of course, forcing partitions to be non-overlapping restricts implementation freedom. This subsection proposes a way to allow more sophisticated partitionings.

[2] The UDF in this case is a nice example of *semantic compression*, where domain-specific information is used to generate smaller representations than a content-blind algorithm would produce.

The core problem partitioning must solve is how to bind state to inputs: i.e., given a modification to a piece of meta data f's state, we must test all inputs that depend on that state. State partitioning does this by grouping all inputs to owns within a single partition. A more sophisticated alternative is to allow overlap between partitions. Naively, we might expect that allowing overlap requires that we have a function enum that, given a piece of state, enumerates all inputs that depend on it. The need for such a function would severely restrict the data structures we could verify, since most do not naturally support the ability to enumerate all inputs given an arbitrary byte of state. However, if we exploit the libFS, we can have it give us the inputs that could be effected by a state modification and yet be able to test that it does so correctly.

We do this using reference counting: if we can reliably determine how many inputs use a piece of state, we can trust the libFS to find which inputs are dependent. Reference counting enables this split because, given the number of inputs that use a piece of state and a libFS-supplied set of inputs, we can verify that the set is sufficient by invoking owns on each element and verifying that it indeed touches s. If it does, then we know we have all inputs that depend on s, otherwise we return an error. Reference counts are computed by the libFS using a refcnt UDF, which associates each byte of data with a reference count. Note that a correct libFS already tracks roughly the information that refcnt needs in order to implement memory management. I.e., it cannot free memory that is still needed by inputs to f. Thus, it does not seem that the requirement that the libFS can write refcnt restricts its freedom in any real way. Lack of space prevents further discussion; interested readers can refer to [30].

4.2 Overview of XN

Designing a flexible exokernel stable storage system has proven difficult: XN is our fourth design. This section provides an overview of how we use UDFs, the cornerstone of XN; the following sections describe some earlier approaches (and why they failed), and aspects of XN in greater depth.

XN provides access to stable storage at the level of disk blocks, exporting a buffer cache registry (Section 4.4) as well as free maps and other on-disk structures. The main purpose of XN is to determine the access rights of a given principal to a given disk block as efficiently as possible. XN must prevent a malicious user from claiming another user's disk blocks as part of her own files. On a conventional OS, this task is easy, since the kernel itself knows the file's meta data format. On an exokernel, where files have application-defined meta data layouts, the task is more difficult. On XN libFSes provide UDFs that act as meta data translation functions specific to each file type. XN uses UDFs to analyze meta data and translate it into a simple form the kernel understands. A libFS developer can install UDFs to introduce new on-disk meta data formats. UDFs allow the kernel to safely and efficiently handle any meta data layout without understanding the layout itself.

UDFs are stored on disk in structures called *templates*. Each template corresponds to a particular meta data format; for example, a UNIX file system would have templates for data blocks, inode blocks, inodes, indirect blocks, etc. Each template T has two UDFs: *owns-udf*$_T$ and *refcnt-udf*$_T$, and two untrusted and potentially nondeterministic functions: *acl-uf*$_T$ and *size-uf*$_T$. All four functions are specified in the same language but only *owns-udf*$_T$ and *refcnt-udf*$_T$ must be deterministic. The other two can have access to, for example, the time of day. The limited language used to write these functions is a pseudo-RISC assembly language, checked by the kernel to ensure determinacy. Once a template is specified, it cannot be changed.

The *owns-udf* function allows XN to check the correctness of libFS meta data modifications (specified as a list of byte range modifications) using techniques from the previous subsection.

The *refcnt-udf* function allows libFSes to represent reference counts however they wish. For simplicity, reference count are stored in the block that is pointed to (and, thus, they are persistent). Whenever an edge to this block is formed, XN verifies that *refcnt-udf* increases by one. And, when an edge is deleted, that the count is decremented.

The *acl-uf* function implements template-specific access control and semantics; its input is a piece of meta data, a proposed modification to that meta data, and set of credentials (e.g., capabilities). Its output is a Boolean value approving or disapproving of the modification. XN runs the proper *acl-uf* function before any meta data modification. *acl-uf*s can implement access control lists, as well as providing certain other guarantees; for example, an *acl-uf* could ensure that inode modification times are kept current by rejecting any meta data changes that do not update them.

The *size-uf* function simply returns the size of a data structure in bytes.

4.3 XN: Problem and history

The most difficult requirement for XN is efficiently determining the access rights of a given principal to a given disk block. We discuss the successive approaches that we have pursued.

Disk-block-level multiplexing. One approach is to associate with each block or extent a capability (or access control list) that guards it. Unfortunately, if the capability is spatially separated from the disk block (e.g., stored separately in a table), accessing a block can require two disk accesses (one to fetch the capability and one to fetch the block). While caching can mitigate this problem to a degree, we are nervous about its overhead on disk-intensive workloads. An alternative approach is to co-locate capabilities with disk blocks by placing them immediately before a disk block's data [54]. Unfortunately, on common hardware, reserving space for a capability would prevent blocks from being multiples of the page size, adding overhead and complexity to disk operations.

Self-descriptive meta data. Our first serious attempt at efficient disk multiplexing provided a means for each instance of meta data to describe itself. For example, a disk block would start with some number of bytes of application-specific data and then say "the next ten integers are disk block pointers." The complexity of space-efficient self-description caused us to limit what meta data could be described. We discovered that this approach both caused unacceptable amounts of space overhead and required excessive effort to modify existing file system code, because it was difficult to shoehorn existing file system data structures into a universal format.

Template-based description. Self-description and its problems were eliminated by the insight that each file system is built from only a handful of different on-disk data structures, each of which can be considered a type. Since the number of types is small, it is feasible to describe each type only once per file system—rather than once per instance of a type—using a *template*.

Originally, templates were written in a declarative description language (similar to that used in self-descriptive meta data) rather than UDFs. This system was simple and better than self-descriptive meta data, but still exhibited what we have come to appreciate as an indication that applications do not have enough control: the system made too many tradeoffs. We had to make a myriad of decisions about which base types were available and how they were represented (how large disk block pointers could be, how the type layout could change, how extents were specified). Given the variety of on-disk data structures described in the file system literature, it seems unlikely that any fixed set of components will ever be enough to describe all useful meta data.

Our current solution uses templates, but trades the declarative description language for a more expressive, interpreted language—UDFs. This lets libFSes track their own access rights without XN understanding how they do so; XN merely verifies that libFSes track block ownership correctly.

4.4 XN: Design and implementation

We first describe the requirements for XN and then present the design.

Requirements and approach

In our experience so far, the following requirements have been sufficient to reconcile application control with protected sharing.

1. To prevent unauthorized access, every operation on disk data must be guarded. For speed, XN uses *secure bindings* [25] to move access checks to bind time rather than checking at every access. For example, the permission to read a cached disk block is checked when the page is inserted into the page table of the libFS's environment, rather than on every access.

2. XN must be able to determine unambiguously what access rights a principal has to a given disk block. For speed, it uses the UDF mechanism to protect disk blocks using the libFS's own meta data rather than guarding each block individually.

3. XN must guarantee that disk updates are ordered such that a crash will not incorrectly grant a libFS access to data it either has freed or has not allocated. This requirement means that meta data that is persistent across crashes cannot be written when it contains pointers to uninitialized meta data, and that reallocation of a freed block must be delayed until all persistent pointers to it have been removed.

While isolation allows separate libFSes to coexist safely, protected sharing of file system state by mutually distrustful libFSes requires three additional features:

1. Coherent caching of disk blocks. Distributed, per-application disk block caches create a consistency problem: if two applications obliviously cache the same disk block in two different physical pages, then modifications will not be shared. XN solves this problem with an in-kernel, system-wide, protected cache registry that maps cached disk blocks to the physical pages holding them.

2. Atomic meta data updates. Many file system updates have multiple steps. To ensure that shared state always ends up in a consistent and correct state, libFSes can lock cache registry entries. (Future work will explore optimistic concurrency control based on versioning.)

3. Well-formed updates. File abstractions above the XN interface may require that meta data modifications satisfy invariants (e.g., that link counts in inodes match the number of associated directory entries). UDFs allow XN to guarantee such invariants in a file-system-specific manner, allowing mutually distrustful applications to safely share meta data.

XN controls only what is necessary to enforce these protection rules. All other abilities—I/O initiation, disk block layout and allocation policies, recovery semantics, and consistency guarantees—are left to untrusted libFSes.

Ordered disk writes

Another difficulty XN must face is guaranteeing the rules Ganger and Patt [36] give for achieving strict file system integrity across crashes: First, never reuse an on-disk resource before nullifying all previous pointers to it. Second, never create persistent pointers to structures before they are initialized. Third, when moving an on-disk resource, never reset the old pointer in persistent storage before the new one has been set.

The first two rules are required for global system integrity—and thus must be enforced by XN—while a file system violating the third rule will only affect itself.

The rules are simple but difficult to enforce efficiently: a naive implementation will incur frequent costly synchronous disk writes. XN allows libFSes to address this by enforcing the rules without legislating how to follow them. In particular, libFSes can choose any operation order which satisfies the constraints.

The first rule is implemented by deferring a block's deallocation until all on-disk pointers to that block have been deleted; a reference count performed at crash recovery time helps libFSes implement the third rule.

The second rule is the hardest of the three. To implement it, XN keeps track of *tainted* blocks. Any block is considered tainted if it points either to an uninitialized block or to a tainted block. LibFSes must not be allowed to write a tainted block to disk. However, two exceptions allow XN to enforce the general rule more efficiently:

First, XN allows entire file systems to be marked “temporary” (i.e., not persistent across reboots). Since these file systems are not persistent, they are not required to adhere to any of the integrity rules. This technique allows memory-based file systems to be implemented with no loss of efficiency.

The second exception is based on the observation that unattached subtrees—trees whose root is not reachable from any persistent root—will not be preserved across reboots and thus, like temporary trees, are free of any ordering constraints. Thus, XN does not track tainted blocks in an unreachable tree until it is connected to a persistent root.

The buffer cache registry

Finally, we discuss the XN buffer cache registry, which allows protected sharing of disk blocks among libFSes. The registry tracks the mapping of cached disk blocks and their meta data to physical pages (and vice versa). Unlike traditional buffer caches, it only records the mapping, not the disk blocks themselves. Because the registry exists independently of libFSes, it allows cached blocks to persist across process invocations. The disk blocks are stored in application-managed physical-memory pages. The registry tracks both the mapping and its state (dirty, out of core, uninitialized, locked). To allow libFSes to see which disk blocks are cached, the buffer cache registry is mapped read-only into application space.

Access control is performed when a libFS attempts to map a physical page containing a disk block into its address space, rather than when that block is requested from disk. That is, registry entries can be inserted without requiring that the object they describe be in memory. Blocks can also be installed in the registry before their template or parent is known. As a result, libFSes have significant freedom to prefetch.

Registry entries are installed in two ways. First, an application that has write access to a block can directly install a mapping to it into the registry. Second, applications that do not have write access to a block can indirectly install an entry for it by performing a "read and insert," which tells the kernel to read a disk block, associate it with an application-provided physical page, set the protection of that page page appropriately, and insert this mapping into the registry. This latter mechanism is used to prevent applications that do not have permission to write a block from modifying it by installing a bogus in-core copy.

XN does not replace physical pages from the registry (except for those freed by applications), allowing applications to determine the most appropriate caching policy. Because applications also manage virtual memory paging, the partitioning of disk cache and virtual memory backing store is under application control. To simplify the application's task and because it is inexpensive to provide, XN maintains an LRU list of unused but valid buffers. By default, when LibOSes need pages and none are free, they recycle the oldest buffer on this LRU list.

XN allows any process to write "unowned" dirty blocks to disk (i.e., blocks not associated with a running process), even if that process does not have write permission for the dirty blocks. This allows the construction of daemons that asynchronously write dirty blocks. LibFSes do not have to trust daemons with write access to their files, only to flush the blocks. This ability has three benefits. First, the contents of the registry can be safely retained across process invocations rather than having to be brought in and paged out on creation and exit. Second, this design simplifies the implementations of libFSes, since a libFS can rely on a daemon of its choice to flush dirty blocks even in difficult situations (e.g., if the application containing the libFS is swapped out). Third, this design allows different write-back policies.

4.5 XN usage

To illustrate how XN is used, we sketch how a libFS can implement common file system operations. These two setup operations are used to install a libFS:

Type creation. The libFS describes its types by storing templates, described above in Section 4.2, into a *type catalogue.* Each template is identified by a unique string (e.g., "FFS Inode"). Once installed, types are persistent across reboots.

LibFS persistence. To ensure that libFS data is persistent across reboots, a libFS can register the root of its tree in XN's *root catalogue.* A root entry consists of a disk extent and corresponding template type, identified by a unique string (e.g., "mylibFS").

After a crash, XN uses these roots to garbage-collect the disk by reconstructing the free map. It does so by logically traversing all roots and all blocks reachable from them: it marks reachable blocks as allocated, non-reachable blocks as free. If this step is too expensive, it could be eliminated by ordering writes to the free map (so that the map always held a conservative picture of what blocks were free) or using a log.

During reconstruction XN also checks for errors in meta data reference counts by counting all pointers to all meta data instances. If this count does not match a meta data's reference count, XN records this violation in an error log. After finding all errors, it then runs libFS-supplied patch programs to fix these and any libFS-specific errors. If errors remain after this process, XN marks the root of any tree that contains a bogus reference count as "tainted." These errors must be fixed before the tree can be used. We discuss reconstruction further in the next section.

After initialization, the new libFS can use XN. We describe a simplified version of the most common operations.

Startup. To start using XN, a libFS loads its root(s) and any types it needs from the root catalogue into the buffer cache registry. Usually both will already be cached.

Read. Reading a block from disk is a two-stage process, where the stages can be combined or separated. First, the libFS creates entries in the registry by passing block addresses for the requested disk blocks and the meta data blocks controlling them (their *parents*). The parents must already exist in the registry—libFSes are responsible for loading them. XN uses *owns-udf* to determine if the requested blocks are controlled by the supplied meta data blocks and, if so, installs registry entries.

In the second stage, the libFS initiates a read request, optionally supplying pages to place the data in. Access control through *acl-uf* is performed at the parent (e.g., if the data loaded is a bare disk block), at the child (e.g., if the data is an inode), or both.

A libFS can load any block in its tree by traversing from its root entry, or optionally by starting from any intermediate node cached in the registry. Note that XN specifically disallows meta data blocks from being mapped read/write.

To speculatively read a block before its parent is known, a libFS can issue a raw read command. If the block is not in the registry, it will be marked as "unknown type" and a disk request initiated. The block cannot be used until after it is bound to a parent by the first stage of the read process, which will determine its type and allow access control to be performed.

Allocate. A libFS selects blocks to allocate by reading XN's map of free blocks, allowing libFSes to control file layout and grouping. Free blocks are allocated to a given meta data node by calling XN with the meta data node, the blocks to allocate, and the proposed modification to the meta data node. XN checks that the requested blocks are free, runs the appropriate *acl-uf* to see if the libFS has permission to allocate, and runs *owns-udf*, as described in Section 4.2, to see that the correct block is being allocated. If these checks all succeed, the meta data is changed, the allocated blocks are removed from the free list, and any allocated meta data blocks are marked tainted (see Section 4.4).

Write. A libFS writes dirty blocks to disk by passing the blocks to write to XN. If the blocks are not in memory, or they have been pinned in memory by some other application, the write is prevented. The write also fails if any of the blocks are tainted and reachable from a persistent root. Otherwise, the write succeeds. If the block was previously tainted and now is not (either by eliminating pointers to uninitialized meta data or by becoming initialized itself), XN modifies its state and removes it from the tainted list.

Since applications control what is fetched and what is paged out when (and in what order), they can control many disk management policies and can enforce strong stability guarantees.

Deallocate. XN uses UDFs to check deallocate operations analogously to allocate operations. If there are no on-disk pointers to a deallocated disk block, XN places it on the free list. Otherwise, XN enqueues the block on a "will free" list until the block's reference count is zero. Reference counts are decremented when a parent that had an on-disk pointer to the block deletes that pointer via a write.

4.6 Crash Recovery Issues

This section discusses two issues of reconstruction: the difficulty of resolving file system errors due to XN's lack of understanding of libFS semantics, and providing more flexible mechanisms for guaranteeing invariants in the presence of crashes.

As we discussed above, XN marks file system roots that contain meta data with erroneous reference counts as "tainted." An apparently simpler solution would be to fix the counts directly. There are two reasons XN does not do so. First, it does not understand libFS meta data syntax. While it could use libFS "helper" functions to fix reference counts, it cannot rely on being able to do so, since they could contain errors. Second, while XN could perhaps use UDF-style techniques to verify helper functions, the action to take in the face of a violation depends on libFS semantics. For example, consider the case of a erroneously high reference count. It is not clear how to correct this count, since it could have arisen in multiple ways. The meta data could have been being moved from one directory to another by a libFS that conservatively deletes the old "source" edge only after the new "destination" edge has been written to disk. (In which case the old edge would have to be removed.) Or an edge could have been removed before the reference count was persistently decremented. (In which case the reference count should be decremented.) XN takes the view that resolving these ambiguities is best left to libFSes.

Guaranteeing that file system invariants hold across system crashes has been an active area of file system research. However, other than control over write orders, XN provides little flexibility in this area. We discuss improvements below. A simple way to improve XN would be to incorporate the notion of logging into it. Violations and (possibly) libFS annotations could be written into a "write-ahead" log that was written to disk before persistent state was updated. In this case, any on disk violations can be removed by simply performing the log actions. To increase flexibility, the format of log records could be controlled by a UDF.

This approach works, and does not appear too hard to implement. Unfortunately, it has serious problems. First, it requires that XN pre-empt many design decisions about how to implement logging. Second, and more seriously, it is not robust: it requires that we anticipate logging and put support for it in XN rather than having logging "fall out" of a general recovery mechanism in XN. The main game of exokernels is extensibility — that all uses of a system can be implemented using its interface — having to anticipate a particular extension (logging) shows its interface is weak.

Fortunately, it appears that we can use UDF techniques to allow libFSes to implement recovery in a completely general way. Consider what XN really needs for valid crash recovery: all it must guarantee is that if a libFS violates an invariant that XN cares about, that on reboot either (1) the violation goes away (because it only affected volatile state) or (2) that the libFS tells XN about it. A log is one way to get the latter guarantee, there are others. One way is to

have the libFS signal violations: rather than forcing libFSes to write blocks out in certain order to prevent violations or using a log to track them persistently, XN has libFSes provide a "reboot UDF" (reboot) that, when given a file system tree walks down the tree emitting any violations. How it detects violations is not XN's concern, it need only verify that reboot will. XN performs this verification as follows by first recording all violations in volatile memory and then, when a libFS wants to do a write that would violate some invariant (e.g., a stably writing a pointer to an uninitialized piece of meta data), XN checks that the associated libFS' reboot UDF would inform it about this violation. XN does so in a way similar to owns verification by (conceptually) running the reboot UDF on the pre-write stable state, performing the disk write, then running the reboot UDF on the post-write stable state. The obvious problem with this approach is efficiency. It is obviously not practical to examine the entire disk on every disk write. Fortunately, it appears that by using both a variation of state partitioning (based on continuation passing [29]) and checkable hints from the file system, we can checkpoint the UDF precisely, to the point that verification need only examine a few bytes in one or two cached disk blocks. We are currently designing this scheme.

4.7 C-FFS: a library file system

This subsection briefly describes C-FFS (co-locating fast file system [37])—a UNIX-like library file system we built—with special reference to additional protection guarantees it provides.

XN provides the basic protection guarantees needed for file system integrity — that both meta data block pointers and reference counts are correct, and that block pointers are correctly typed. But real-world file systems often require other, file-system-specific invariants. For instance, UNIX file systems must ensure the uniqueness of file names within a directory. This type of guarantee can be provided in any number of ways: in the kernel, in a server, or, in some cases, by simple defensive programming. C-FFS currently downloads methods into the kernel to check its invariants. We are currently developing a system similar to UDFs that can be used to enforce type-specific invariants in an efficient, extensible way.

Our experience with C-FFS shows that, even with the strongest desired guarantees, a protected interface can still provide significant flexibility to unprivileged software, and that the exokernel approach can deal as readily with high-level protection requirements as it can with those closer to hardware.

C-FFS makes four main additions to XN's protection mechanisms:

1. Access control: it maps the UNIX representation and semantics of access control (uids and gids, etc.) to those of exokernel capabilities.

2. Well-formed updates: C-FFS guarantees UNIX-specific file semantics: for example, that directories contain legal, aligned file names.

3. Atomicity: C-FFS performs locking to ensure that its data is always recoverable and disk writes only occur when meta data is internally consistent.

4. Implicit updates: C-FFS ensures that certain state transitions are implicit on certain actions. Some examples are that modification times are updated when file data are changed, and that renaming or deleting a file updates the name cache.

It is not difficult to implement UNIX protection without significantly degrading application power. C-FFS protection is implemented mainly by a small number of if-statements rather than by procedures that limit flexibility. The most intricate operation—ensuring that files in a directory have unique names—is less than 100 lines of code that scans through a linked list of cached directory blocks to ensure name uniqueness.

4.8 Discussion

Similar to how XN uses UDFs to interpret libFS meta data, libFSes can use UDFs to traverse other file system's meta, even when they do not understand its syntax. (Of course, modification of this meta data typically requires such understanding.)

UDFs add a negligible cost to XN. As the next section discusses, we have run file system benchmarks with and without XN enabled and its overhead (and, thus, the overhead of UDFs) is lost in experimental noise. There are two reasons for this. First, protection tends to be off of the critical path. Second, XN operations tend to be embedded

in heavy-weight disk operations. For example, most block operations on cached blocks can simply interact with the buffer cache registry, rather than using to UDFs. Even if neither of these two conditions held, UDFs would not add significant overhead: most UDFs tend to be fairly simple, and if they were directly executed rather than interpreted, cost a few tens of instructions.

XN allows "nestable" extension of file systems. Because XN verifies the correctness of reference counts, pointers, and meta data interpreters, it allows untrusted implementors to extend an existing file system without compromising its integrity. Thus, it is possible to add an entirely new directory type to a file system and have it point to old types, perform access control on them, etc. without the existing implementation open to malice. We do not know of any other way to achieve this same result.

Stable storage is the most challenging resource we have multiplexed. Future work will focus on two areas. First, we plan to implement a range of file systems (log-structured file systems, RAID, and memory-based file systems), thus testing if the XN interface is powerful enough to support concurrent use by radically different file systems. Second we will investigate using lightweight protected methods like UDFs to implement the simple protection checks required by higher-level abstractions.

Chapter 5

Performance of exokernel systems

> The fundamental principle of science, the definition almost, is this: the sole test of the validity of any idea is experiment. – Richard P. Feynman
>
> Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that! Charles Lutwidge Dodgson (Lewis Carroll) (1832-1898) "Through the Looking Glass"

This chapter focuses on four questions:

1. Do common, unaltered applications benefit from an exokernel? To answer this question we present performance results of Unix applications on Xok that have been linked against an optimized libFS. Applications on Xok run comparably or significantly faster compared to both FreeBSD and OpenBSD.

2. Is exokernel flexibility costly? To partially answer this we present performance numbers of the previous experiment running on top of a re-implementation of the libFS, the C-FFS file system, within OpenBSD. The exokernel and OpenBSD perform roughly comparably.

3. Are aggressive applications significantly times faster? Among other experiments, we compare the performance of an optimized webserver, which is eight times faster than on traditional systems.

4. Does local control lead to bad global performance? An exokernel gives applications significantly more control than traditional operating systems do. Can it reconcile strong local control with good global performance?

   To answer this question we measure aggressive workloads on Xok and FreeBSD: (1) given the same workload, an exokernel performs comparably to widely used monolithic systems, and (2) when local optimizations are performed, that whole system performance improves, sometimes significantly.

We describe our experimental environment below. The remainder of the chapter addresses each question in turn.

5.1 Xok Experimental Environment

We compare Xok/ExOS to both FreeBSD 2.2.2 and OpenBSD 2.1 on the same hardware. Xok uses device drivers that are derived from those of OpenBSD. ExOS also shares a large source code base with OpenBSD, including most applications and most of libc. Compared to OpenBSD and FreeBSD, ExOS has not had much time to mature; we built the system in less than two years and moved to the x86 platform only a year ago.

All experiments are performed on 200-MHz Intel Pentium Pro processors with a 256-KByte on-chip L2 cache and 64-MByte of main memory. The disk system consists of an NCR 815 SCSI controller connecting a fast SCSI chain with one or more Quantum Atlas XP32150 disk drives on the PCI bus (vs440fx PCI chip set). Reported times are the minimum time of ten trials (the standard deviations of the total run times are less than three percent).

As discussed in Chapter 3, some ExOS data structures do not have full protection. To compensate for this lack, we have artificially inserted three extra system calls before every write to shared tables. This gives a pessimal feel for the

Benchmark	Description (application)
Copy small file	copy the compressed archived source tree (cp)
Uncompress	uncompress the archive (gunzip)
Copy large file	copy the uncompressed archive (cp)
Unpack file	unpack archive (pax)
Copy large tree	recursively copy the created directories (cp).
Diff large tree	compute the difference between the trees (diff)
Compile	compile source code (gcc)
Delete files	delete binary files (rm)
Pack tree	archive the tree (pax)
Compress	compress the archive tree (gzip)
Delete	delete the created source tree (rm)

Table 5.1: The I/O-intensive workload installs a large application (the lcc compiler). The size of the compressed archive file for lcc is 1.1 MByte.

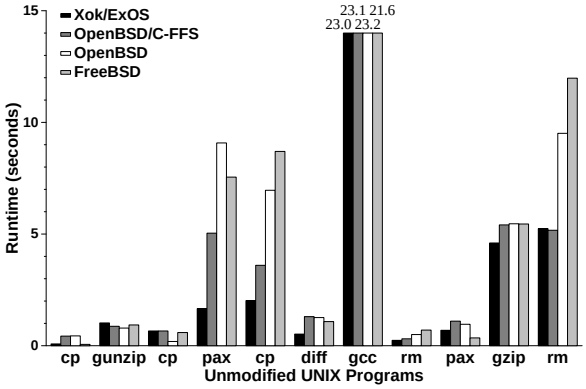

Figure 5-1: Performance of unmodified UNIX applications. Xok/ExOS and OpenBSD/C-FFS use a C-FFS file system while Free/OpenBSD use their native FFS file systems. Times are in seconds.

cost of protection. All measurements reported in this thesis include these extra calls. In practice, protection is off the critical path, and these overheads are lost in experimental noise.

It is important to note that a sufficiently motivated kernel programmer can implement any optimization that is implemented in an extensible system. In fact, a member of our research group, Costa Sapuntzakis, has implemented a version of C-FFS within OpenBSD. Extensible systems (and we believe exokernels in particular) make these optimizations significantly easier to implement than centralized systems do. For example, porting C-FFS to OpenBSD took more effort than designing C-FFS and implementing it as a library file system. The experiments below demonstrate that by using unprivileged application-level resource management, any skilled programmer can implement useful OS optimizations. The extra layer of protection required to make this application-level management safe costs little.

5.2 Performance of common, unaltered applications

If an exokernel only improves the performance of strange, niche applications or, requires that applications be modified to benefit, then its usefulness is severely diminished.

Our experiments show that an exokernel matters, even for common applications. We measure the performance of an I/O-intensive software development workload made up of the mainstream applications listed in Table 5.1 (most are found in "/usr/bin" on a Unix system). As Figure 5-1 shows, by linking against an optimized library file system (C-FFS), some unaltered UNIX applications run significantly faster on top of Xok/ExOS than identical versions executed on traditional systems. Xok/ExOS completes all benchmarks in 41 seconds, 19 seconds faster than FreeBSD and OpenBSD. On eight of the eleven benchmarks Xok/ExOS performs better than Free/OpenBSD (in one case by over a factor of four). ExOS's performance improvements are due to its C-FFS file system.

In general, normal applications benefit from an exokernel by being linked against an optimized libOS. Thus, even

without modifications, they still profit from an exokernel. While an exokernel allows experimentation with completely different OS interfaces, a more important result may be improving the rate of innovation of implementations of existent interfaces.

We also ran the Modified Andrew Benchmark (MAB) [69]. On this benchmark, Xok/ExOS takes 11.5 seconds, OpenBSD/C-FFS takes 12.5 seconds, OpenBSD takes 14.2 seconds, and FreeBSD takes 11.5 seconds. The difference in performance on MAB is less profound than on the I/O-intensive benchmark, because MAB stresses fork, an expensive function in Xok/ExOS. ExOS's fork performance suffers because Xok does not yet allow environments to share page tables. Fork takes six milliseconds on ExOS, compared to less than one millisecond on OpenBSD.

5.3 The cost of exokernel flexibility

An exokernel provides extreme flexibility. Rightfully, any systems builder would expect that this flexibility would have a performance cost. This section looks at two possible costs: (1) the opportunity cost of OS abstractions at library level, and (2) the brute cost of protection,

5.3.1 The cost of OS abstractions in libraries

Given the same application code and same OS code (placed in a libOS on an exokernel, in the kernel on a monolithic system), does a traditional system would provide superior performance? Or, phrased another way, if an optimization is done on an exokernel and gives a factor of four, would the same optimization done on a traditional OS give even more?

To partially answer this question we measure the performance of an OpenBSD-based implementation of C-FFS (done by Costa Sapuntzakis) on the workload in the previous experiment.

Figure 5-1 shows the performance of these applications over Xok/ExOS and OpenBSD/C-FFS. The total running time for Xok/ExOS is 41 seconds and for OpenBSD/C-FFS is 51 seconds. Since ExOS and OpenBSD/C-FFS use the same type of file system, one would expect that ExOS and OpenBSD perform equally well. As can be seen in Figure 5-1, Xok/ExOS performance is indeed comparable to OpenBSD/C-FFS on eight of the 11 applications. On three applications (pax, cp, diff), Xok/ExOS runs considerably faster (though we do not yet have a good explanation for this).

From these measurements we conclude that, even though ExOS implements the bulk of the operating system at the application level, common software development operations on Xok/ExOS perform comparably to OpenBSD/C-FFS. They demonstrate that—at least for this common domain of applications—an exokernel's flexibility can be provided for free: even without aggressive optimizations ExOS's performance is comparable to that of mature monolithic systems. The cost of low-level multiplexing is negligible: an optimization done on an exokernel can give the same performance improvement as done on a traditional system.

5.3.2 The cost of protection

While the above experiment by no means “proves” that the exokernel structure will never penalize applications, it agrees with our experiences that exokernel flexibility does not impose overheads. The main reason for this is that protection tends to be off the critical path. For example, when the benchmarks in in Figure 5-1 are run without XN or any of the extra system calls, the performance difference is “noise”: 39.7 seconds to 41.1 seconds, despite reducing the overall number of Xok system calls from 300,000 to 81,000. Real workloads are dominated by costs other than protection and system call overhead.

Some secondary, more specific reasons for the lack of impact of an extra protection layer are that: in many cases exokernel protection is not a duplication, since it allows library operating systems to remove their corresponding checks. Furthermore, the cost of checking protection (usually a table lookup to map principals to access rights) tends to be dwarfed by the cost of the operations it is coupled to. For example, system calls are roughly an order of magnitude more expensive while I/O operations such as disk or network accesses can be over a 1000 times more costly. In those rare cases where protection checks are more elaborate, they tend to be naturally cachable. In the case of file blocks, for instance, access rights can be stored in the buffer cache along with the block they correspond too.

While an exokernel can impose extra protection checks, it also makes some operations cheaper: using a library operating system means that many operations that formerly required system calls now can be performed with function

calls. Thus, for some uses an exokernel is intrinsically faster than a more traditional system. [1] For example, an OpenBSD emulator written on our most recent exokernel, Xok, sometimes runs OpenBSD application binaries faster than their non-emulated performance on OpenBSD for the simple reason that that many system calls (e.g., to read data structures) become function calls into the emulator's libOS.

5.4 Aggressive application performance

In part, the exokernel architecture was motivated by the ability to perform application-specific or, more commonly, domain-specific optimizations. [2] Two natural questions, then, are first, does an exokernel give sufficient power that interesting optimizations can be done? Second, do domain-specific optimizations yield significant improvements or do they only give "noise" level improvements?

Experiments show that an exokernel enables domain-specific optimizations that give order-of-magnitude performance improvements [48]. Our specific results come from an interface designed for fast I/O, which improves the performance of applications such as web servers.

5.4.1 XCP: a "zero-touch" file copying program

XCP is an efficient file copy program. It exploits the low-level disk interface by removing artificial ordering constraints, by improving disk scheduling through large schedules, by eliminating data touching by the CPU, and by performing all disk operations asynchronously.

Given a list of files, XCP works as follows. First, it enumerates and sorts the disk blocks of all files and issues large, asynchronous disk reads using this schedule. (If multiple instances of XCP run concurrently, the disk driver will merge the schedules.) Second, it creates new files of the correct size, overlapping inode and disk block allocation with the disk reads. Finally, as the disk reads complete, it constructs large writes to the new disk blocks using the buffer cache entries. This strategy eliminates all copies; the file is DMAed into and out of the buffer cache by the disk controller—the CPU never touches the data.

XCP is a factor of three faster than the copy program (CP) on Xok/ExOS that uses UNIX interfaces, irrespective of whether all files are in core (because XCP does not touch the data) or on disk (because XCP issues disk schedules with a minimum number of seeks and the largest contiguous ranges of disk blocks).

The fact that the file system is an application library allows us both to have integration when appropriate and to craft new abstractions as needed. This latter ability is especially profitable for the disk both because of the high cost of disk operations and because of the demonstrated reluctance of operating systems vendors to provide useful, simple improvements to their interfaces (e.g., prefetching, asynchronous reads and writes, fine-grained disk restructuring and "sync" operations).

5.4.2 The Cheetah HTTP/1.0 Server

The exokernel architecture is well suited to building fast servers (e.g., for NFS servers or web servers). Server performance is crucial to client/server applications [45], and the I/O-centric nature of servers makes operating system-based optimizations profitable.

Greg Ganger has developed an extensible I/O library (XIO) for fast servers and a sample application that uses it, the Cheetah HTTP server. This library is designed to allow application writers to exploit domain-specific knowledge and to simplify the construction of high-performance servers by removing the need to "trick" the operating system into doing what the application requires (e.g., Harvest [14] stores cached pages in multiple directories to achieve fast name lookup).

An HTTP server's task is simple: given a client request, it finds the appropriate document and sends it. The Cheetah Web server performs the following set of optimizations as well as others not listed here.

Merged File Cache and Retransmission Pool. Cheetah avoids all in-memory data touching (by the CPU) and the need for a distinct TCP retransmission pool by transmitting file data directly from the file cache using precomputed

[1] Admittedly, in our experience, even high system call overhead has little impact on real application performance. But this point argues even further for the non-issue of the cost of having operating system code in libraries.

[2] In fact, the original exokernel paper [25] focuses almost exclusively on application-specific optimizations rather improving the rate of whole-system innovations, which we have come to regard as a more significant benefit.

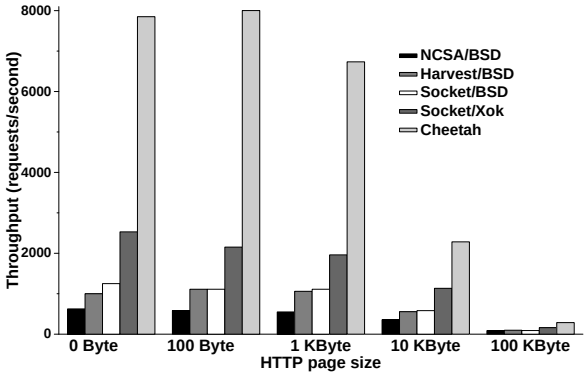

Figure 5-2: HTTP document throughput as a function of the document size for several HTTP/1.0 servers. **NCSA/BSD** represents the NCSA/1.4.2 server running on OpenBSD. **Harvest/BSD** represents the Harvest proxy cache running on OpenBSD. **Socket/BSD** represents our HTTP server using TCP sockets on OpenBSD. **Socket/Xok** represents our HTTP server using the TCP socket interface built on our extensible TCP/IP implementation on the Xok exokernel. **Cheetah/Xok** represents the Cheetah HTTP server, which exploits the TCP and file system implementations for speed.

file checksums (which are stored with each file). Data are transmitted (and retransmitted, if necessary) to the client directly from the file cache without CPU copy operations. (Cao et al. have also used this technique [70].)

Knowledge-based Packet Merging. Cheetah exploits knowledge of its per-request state transitions to reduce the number of I/O actions it initiates. For example, it avoids sending redundant control packets by delaying ACKs on client HTTP requests, since it knows it will be able to piggy-back them on the response. This optimization is particularly valuable for small document sizes, where the reduction represents a substantial fraction (e.g., 20%) of the total number of packets.

HTML-based File Grouping. Cheetah co-locates files included in an HTML document by allocating them in disk blocks adjacent to that file when possible. When the file cache does not capture the majority of client requests, this extension can improve HTTP throughput by up to a factor of two.

Figure 5-2 shows HTTP request throughput as a function of the requested document size for five servers: the NCSA 1.4.2 server [68] running on OpenBSD 2.0, the Harvest cache [14] running on OpenBSD 2.0, the base socket-based server running on OpenBSD 2.0 (i.e., our HTTP server without any optimizations), the base socket-based server running on the Xok exokernel system (i.e., our HTTP server without any optimizations with vanilla socket and file descriptor implementations layered over XIO), and the Cheetah server running on the Xok exokernel (i.e., our HTTP server with all optimizations enabled).

Figure 5-2 provides several important pieces of information. First, our base HTTP server performs roughly as well as the Harvest cache, which has been shown to outperform many other HTTP server implementations on general-purpose operating systems. Both outperform the NCSA server. This gives us a reasonable starting point for evaluating extensions that improve performance. Second, the default socket and file system implementations built on top of XIO perform significantly better than the OpenBSD implementations of the same interfaces (by 80–100%). The improvement comes mainly from simple (though generally valuable) extensions, such as packet merging, application-level caching of pointers to file cache blocks, and protocol control block reuse.

Third, and most importantly, Cheetah significantly outperforms the servers that use traditional interfaces. By exploiting Xok's extensibility, Cheetah gains a four times performance improvement for small documents (1 KByte and smaller), making it eight times faster than the best performance we could achieve on OpenBSD. Furthermore, the large document performance for Cheetah is limited by the available network bandwidth (three 100Mbit/s Ethernets) rather than by the server hardware. While the socket-based implementation is limited to only 16.5 MByte/s with 100% CPU utilization, Cheetah delivers over 29.3 MByte/s with the CPU idle over 30% of the time. The extensibility of ExOS's default unprivileged TCP/IP and file system implementations made it possible to achieve these performance improvements incrementally and with low complexity.

The optimizations performed by Cheetah are architecture independent. In Aegis, Cheetah obtained similar performance improvements over Ultrix web servers [49].

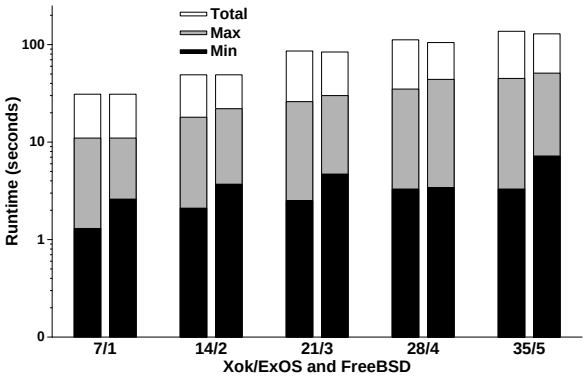

Figure 5-3: Measured global performance of Xok/ExOS (the first bar) and FreeBSD (the second bar), using the first application pool. Times are in seconds and on a log scale. *number/number* refers to the the total number of applications run by the script and the maximum number of jobs run concurrently. **Total** is the total running time of each experiment, **Max** is the longest runtime of any process in a given run (giving the worst latency). **Min** is the minimum.

5.5 Global performance

An exokernel gives applications significantly more control than traditional operating systems do. However, it must also guarantee good global performance when running multiple applications concurrently. The experiments in this section measure the situation where the exokernel architecture seems potentially weak: under substantial load where selfish applications are consuming large resources and utilizing I/O devices heavily. The results indicate that an exokernel can successfully reconcile local control with global performance: (1) given the same workload, an exokernel performs comparably to widely used monolithic systems, and (2) that when local optimizations are performed, that whole system performance improves, and can do so significantly.

There are two intuitions behind these results. First, most local optimizations, because they make applications run faster, lead to more resources globally. For example, if an application, after being linked against an optimized libOS cuts its runtime from ten seconds to one second, then there are nine seconds of freed resources to go around the entire system. Second, an exokernel mediates allocation and revocation of resources. Therefore it has the power to enforce any global policy that a traditional operating system can. Thus, all else equal, it has no problem achieving similar global performance. The single new challenge an exokernel faces is deriving information lost by dislocating abstractions into application space. For example, traditional operating systems manage both virtual memory and file caching. As a result, they can perform global resource management of pages that takes into account the manner in which a page is being used. In contrast, if an exokernel dislocates virtual memory and file buffer management into library operating systems it no longer can make such distinctions. While such information matters in theory, in practice we have found it either unnecessary or crude enough that no special methods have been necessary to derive it. However, whether this happy situation always holds is an open question.

5.5.1 Experiments

Global performance has not been extensively studied. We use the total time to complete a set of concurrent tasks as a measure of system throughput, and the minimum and the maximum latency of individual applications as a measure of interactive performance. For simplicity we compare Xok/ExOS's performance under high load to that of FreeBSD; in these experiments, FreeBSD always performs better than OpenBSD, because of OpenBSD's small, non-unified buffer cache. While this methodology does not guarantee that an exokernel can compare to any centralized system, it does offer a useful relative metric.

The space of possible combinations of applications to run is large. The experiments use randomization to ensure we get a reasonable sample of this space. The inputs are a set of applications to pick from, the total number to run, and the maximum number that can be running concurrently. Each experiment maintains the number of concurrent processes at the specified maximum. The outputs are the total running time, giving throughput, and the time to run each application. Poor interactive performance will show up as a high minimum latency.

The first application pool includes a mix of I/O-intensive and CPU-intensive programs: pack archive (pax -w),

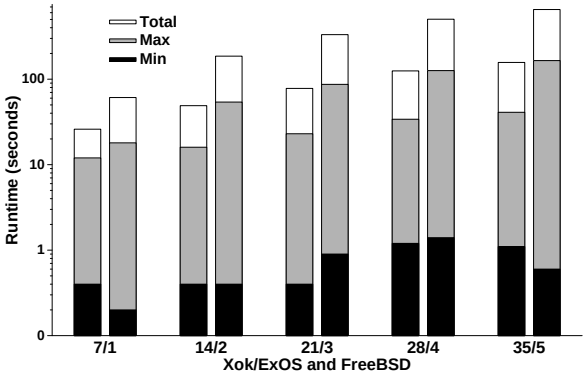

Figure 5-4: Measured global performance of Xok/ExOS (the first bar) and FreeBSD (the second bar), using the second application pool. Methodology and presentation are as described for Figure 5-3.

search for a word in a large file (grep), compute a checksum many times over a small set of files (cksum), solve a traveling salesman problem (tsp), solve iteratively a large discrete Laplace equation using successive overrelaxation (sor), count words (wc), compile (gcc), compress (gzip), and uncompress (gunzip). For this experiment, we chose applications on which both Xok/ExOS and FreeBSD run roughly equivalently. Each application runs for at least several seconds and is run in a separate directory from the others (to avoid cooperative buffer cache reuse). The pseudo-random number generators are identical and start with the same seed, thus producing identical schedules. The applications we chose compete for the CPU, memory, and the disk.

Figure 5-3 shows on a log scale the results for five different experiments: seven jobs with a maximum concurrency of one job through 35 jobs with a maximum concurrency of five jobs. The results show that an exokernel system can achieve performance roughly comparable to UNIX, despite being mostly untuned for global performance.

With a second application pool, we examine global performance when specialized applications (emulated by applications that benefit from C-FFS's performance advantages) compete with each other and non-specialized applications. This pool includes tsp and sor from above, unpack archive (pax -r) from Section 5.2, recursive copy (cp -r) from Section 5.2, and comparison (diff) of two identical 5 MB files. The pax and cp applications represent the specialized applications.

Figure 5-4 shows on a log scale the results for five experiments: seven jobs with a maximum concurrency of one job through 35 jobs with a maximum concurrency of 5 jobs. The results show that global performance on an exokernel system does not degrade even when some applications use resources aggressively. In fact, the relative performance difference between FreeBSD and Xok/ExOS increases with job concurrency.

5.5.2 Discussion

The central challenge in an exokernel system is not *enforcing* a global system policy but, rather, *deriving* the information needed to decide what enforcement involves and doing so in such a way that application flexibility is minimally curtailed. Since an exokernel controls resource allocation and revocation, it has the power to enforce global policies. Quota-based schemes, for instance, can be trivially enforced using only allocation denial and revocation. Fortunately, the crudeness of successful global optimizations allows global schemes to be readily implemented by an exokernel. For example, Xok currently tracks global LRU information that applications can use when deallocating resources.

We believe that an exokernel can provide global performance *superior* to current systems. First, effective local optimization can mean there are more resources for the entire system. Second, an exokernel gives application writers machinery to orchestrate inter-application resource management, allowing them to perform domain-specific global optimizations not possible on current centralized systems (e.g., the UNIX “make” program could be modified to orchestrate the complete build process). Third, an exokernel can unify the many space-partitioned caches in current systems (e.g., the buffer cache, network buffers, etc.). Fourth, since applications can know when resources are scarce, they can make better use of resources when layering abstractions. For example, a web server that caches documents in virtual memory could stop caching documents when its cache does not fit in main memory. Future research will pursue these issues.

5.5.3 Summary

Our experiments show that even common, unaltered applications can benefit on exokernels, simply by being linked against an optimized library operating system. Importantly, libOS optimizations appear just as effective as their equivalent in-kernel implementation. Aggressive applications that want to manage their resources show even greater improvements. The improvement is especially dramatic for I/O-centric applications, such as our web server, which runs up to a factor of 8 faster than its closest equivalent. Finally, the power that an exokernel gives to applications does not lead to poor global performance. In fact, when this control is used to improve application speed, an exokernel system can have dramatically improved global performance, since there are more resources to go around. Based on these experiments, the exokernel architecture appears to be a promising alternative to traditional systems.

Chapter 6

Reflections on Downloading Code

The individual's whole experience is built upon the plan of his language. — Henri Delacroix

Extensibility refers roughly to how easily a system's functionality can be augmented. A strong thread in computer science has been developing techniques to enhance extensibility, ranging from programming methodologies such as structured programming to assist program modification, to dynamic linking of device drivers to add new functions to an operating system kernel. Using language to build extensible systems has had a venerable tradition: the widely-used text editor emacs has done so since its inception [11, 81], database systems exploit it to enrich queries and extend data types, and more recently web browsers and servers have used it to extend their base functionality.

A variety of operating systems have allowed applications to download untrusted code into them as a way to extend their functionality [9, 22, 25, 32, 48, 71, 79, 80, 92]. This chapter documents experiences drawn from the exokernel systems described in this thesis. These experiences cover a period of four years, and span numerous rethinkings of the role of downloaded code, and, as well, much belated realization of its implications and misuses.

The ability to download code has subtle implications. This chapter's central contribution is its perspective on the abilities downloaded code grants and removes, as well as its concrete examples of how these gained and lost abilities matter in practice. Some specific insights include:

1. “Infinite” extensibility requires Turing completeness, Turing completeness gives infinite extensibility.

   Solving the negative problem of extensibility requires supporting all unanticipated uses. A guaranteed solution is to let applications inject general-purpose code into the system, thereby granting them the ability to implement any computable policy or mechanism. (An alternative code motion, uploading operating system code into the application, provides the same guarantee.)

   Conversely, an interface striving for infinite generality is implicitly attempting to provide Turing completeness. Explicitly realizing this fact leads to the obvious solution of having clients pass in general-purpose code.

   The following point provides an example:

2. Correct applications track what resources they have access to, rendering the operating system's bookkeeping redundant. Thus, if the operating system can reuse the application's data structures, it can eliminate this redundancy.

   To ensure that the operating system can understand these structures without restricting their implementation, we use the previous point: clients provide a data structure interpreter, written in a Turing complete language, that the operating system uses to extract the bookkeeping information it needs.

   To force these interpreters to be correct, we use the following technique:

3. Inductive incremental testing provides a practical way to verify the correctness (not just mere safety) of deterministic code. We call functions amenable to this approach *untrusted deterministic functions* (UDFs).

   Using UDFs, operating systems can avoid pre-determining implementation tradeoffs by leaving implementation decisions to client UDFs. Our most interesting use of UDFs, verifying the resource interpreters described above, lets untrusted file systems track what disk blocks they own without the operating system understanding how, yet without being vulnerable to malice.

4. There are practical nuances between using code or data to orchestrate actions between entities. Two examples follow.

   Data is not a Turing machine: compared to downloaded code, data is transparent, and its operations (read and write) trivially bounded in cost and guaranteed to terminate. Code is not, necessarily, any of these things. Injecting potentially non-terminating black boxes into an operating system does not help simplicity. We only use downloaded code in a few specific instances, and have removed it more than once.

   Code requires indirection. Imposing code between the operating system and its data forces it to go through a layer of potentially non-terminating indirection. For example, since our disk subsystem relies on client UDFs to interpret file system structures, it cannot modify them directly, but requires assistance for operations as simple as changing one disk block pointer to another in order to compact the disk.

5. The main benefit of downloaded code is *not* execution speed, but rather trust and consequently power.

   While we started with the view that downloading code was useful for speed (e.g., to eliminate the cost of kernel/user boundary crossings) [25] it has turned out to be far more crucial for power: because downloaded code can be controlled, it can be safely granted abilities that external, unrestricted application code cannot.

   For example, since downloaded file system UDFs can be verified, they can be trusted to track that file system's disk blocks. Obviously, unrestricted applications cannot similarly be trusted.

The chapter is organized as follows. The next four sections discuss different language-based subsystems, which form the spine for our experiences and lessons: DPF [31], our packet filter engine; *application specific message handlers* (ASHs) [25, 92, 93], a networking system which invokes downloaded code on message arrival; XN [48], the disk protection subsystem described in Chapter 4, which contains our most interesting use of downloaded code; and finally, *protected methods*, which applications use to enforce invariants on shared state. Section 6.5 explores some of the slippery implications of using computational building blocks in lieu of passive procedural interfaces. Section 6.6 links our experiences to the existent literature on extensible systems.

6.1 DPF: dynamic packet filters

Packet filters are one of the most successful examples of downloaded code: most modern operating systems provide support for them. This section discusses our packet filter system, DPF ("dynamic packet filters") [25, 31], which was briefly discussed in Chapter 3. DPF uses a combination of language and compilation techniques to get speed and protection. We focus on what abilities downloading code granted, along with some lessons, including our in-hindsight naive mistakes applying dynamic compilation to packet filters.

6.1.1 Language design

The main challenge in DPF is effective detection of filter overlap: i.e., knowing when two filters are comparing the same message offsets to the same values. Protection requires this feature, since otherwise an application could easily steal another's messages. Additionally, overlap detection aids efficiency, since it enables merging of overlapping filter segments [6, 95]. Such merging is crucial for scalability. Typical systems have many filters simultaneously active, one for each network connection.

We made the DPF language declarative, unlike previous packet filter languages [65, 64, 95]. A lower-level imperative language artificially complicates overlap detection due to the presence of operationally different yet functionally equivalent instruction sequences. A declarative language allowed us to avoid puzzling out such artifacts. [1]

A declarative language also assists a sophisticated optimizer, since it makes programmer intent clearer. The small simplicity of the DPF language (e.g., the lack of loops, aliasing, variables) made constructing a sophisticated in-kernel compiler tractable. Given more primitive compilation technology, an imperative language would have been more appropriate, since it would allow a clever programmer to implement optimizations the compiler would not otherwise do.

DPF uses dynamic compilation to compile filters, which makes such semantic incorporation easy. Its optimizations include encoding filter constants in the instruction stream rather than being loaded from a data structure, coalescing

[1] Independently of our work, the PATHFINDER packet filter language was also designed declaratively, though here this form made it easier to put into hardware [6].

comparisons of adjacent message values, estimating alignment of message loads to eliminate the pessimal use of unaligned memory loads and, finally, eliminating bounds checks.

The most unusual optimization DPF does is *hash table compilation*. When filters compare the same message offset to the same value we merge them. However, when they compare to different values we create a hash table holding the constant each filter compares to. DPF uses dynamic code generation to compile this hash table to executable code. For example, if the hash table has no collisions, the lookup can elide collision checks. If there are a small number of keys (say 8 or less) it instead generates a specialized binary search that hard codes each key as an immediate value in the instruction stream. Otherwise it creates a jump table. Additionally, since the number and value of keys are known at runtime, DPF can select among several hash functions to obtain the best distribution, and then encode the chosen function in the instruction stream.

6.1.2 Downloaded code provides power

Because downloaded code can be controlled, it can be safely granted abilities that external, unrestricted application code cannot.

Because packet filters can be restricted (to terminate quickly, and to not overlap) they can be trusted to correctly claim incoming packets. The alternative, asking applications if they want a particular packet, is clearly untenable since there is nothing preventing an application from claiming all packets. The each of the remaining three subsystems we describe illustrate this point.

To restrict code, download it. In general, code on the other side of a trust boundary cannot be prevented from doing arbitrary computations, nor from communicating. Embedding downloaded code in a trusted execution context allows the host to completely control the code, preventing actions otherwise impossible to restrict. DPF overlap detection is one example of restricting computation. Another is letting filters see packets they do not own, since the they cannot leak packet contents to their associated applications. An example outside the context of this chapter is Myers' information flow control work [66], which uses a trusted compiler and runtime environment to prevent applications from leaking sensitive information.

6.1.3 Some lessons

The cost of extensibility. Compared to “hard-wired” systems, extensibility can add noticeable complexity. A static packet demultiplexer is roughly an order of magnitude smaller than DPF's three thousand lines of code. It is simpler as well, consisting of a sequence of conditions and a hash table implementation. The benefits of enabled functionality can make introduction of this layer of indirection worthwhile, but they rarely come without cost.

Code is data, but data is not code. In a theoretical sense, all code is data. However, in a practical sense, some code is more data than others. Data changes frequently, code does not. As a result, it is significantly more difficult to dynamically compile data than code. Data's rate of change typically requires that the implementor use self-modifying code and, worse, understand how to make it fast on modern architectures that penalize the operations it needs.

While a single packet filter code fragment is constant during its life in the kernel, the trie created by merging filters changes incessantly, potentially on each insertion. In a sense, by merging filters we have transmuted them from code (changing rarely) to data (changing frequently). Our first real implementation of DPF missed the implications of this transformation. We naively treated the trie as slowly changing code in that we potentially regenerated the entire structure on insertion of a new filter [31]. This regeneration did not scale at all with large number of branches (which correspond to protocols) in the trie. A second implementation made filter insertion cheaper by incrementally patching the previously generated code. Conceptually, this change was simple: instead of representing each filter operation of “load message value and compare” as a basic block, it became a code fragment linked to the next fragment by an indirect jump. This design isolated the effects of adding a new filter to patching the jumps threading the code together using either self-modifying code or, on machines where instruction cache flushing is expensive, indirect jumps.

Which this modification is conceptually trivial, it conflicts with modern architecture trends. RISC processors require programmers manually implement instruction and data cache coherence via explicit cache flushes. In conditions of frequent insertion (e.g., once every six message matches) this operation adds significant overhead, both in the cost of the flush itself, and in the opportunity cost of subsequent cache misses. [2] There are ways to mitigate this cost, but they complicate the code generator. (For example, generating code into memory guaranteed to not be in the instruction

[2] For all the drawbacks of its instruction set, the x86 family of machines makes self-modifying code simple, since the prevalence in running applications requires that it work, and not cost outrageously.

cache.) Another problem, banal but real, is that threading requires that pointers be loaded as immediates, which on 64-bit machines can be expensive. Again, fixes are not intellectually deep, but can require tedious bookkeeping.

In contrast, dynamically compiling code is much simpler. The code changes infrequently, if ever, and tends to be used many times before being discarded. As a result, compilation is a "one off" affair and its cost easily recouped. While current dynamic compilation techniques works reasonably well for compiling languages, they must mature further before they can be readily used to compile data structures.

6.2 Application-specific message handlers

This section describes lessons learned in the context of the ASH networking subsystem. [25, 92, 93]. ASHs (application-specific message handlers) are application code downloaded into the operating system, made safe using a variant of software fault isolation [89], and invoked upon message arrival.

An exokernel attempts to place most privileged operating system code in the more forgiving and innovative environment of untrusted software. For example, it uses DPF to enable library-based networking. In a sense, ASHs allow the reverse movement in that they allow applications to exchange a more general environment, where they can run unrestricted in time and in operation, for a more restrictive environment that grants two specific abilities: tight coupling to clock and network events, and a way incorporate application semantics into OS actions. These two abilities enable libOSes to efficiently handle incoming messages. Tight coupling to interrupts allows low-latency message replies. Control over over message placement eliminates intermediate buffering and its associated copies. The challenge in this motion comes from forcing ASHs to remain unprivileged (otherwise dynamic linking would suffice).

This section focuses on these two ASH abilities and then closes with experiences, including some mistakes.

6.2.1 Pulling application semantics into event handling

ASHs are cheap to invoke (a few tens of instructions), and their runtime bounded. Thus, they can be run in situations where heavy-weight application scheduling is unacceptable. ASHs thus represent a way for applications to decouple actions from their own execution. Frequently, decoupling happens in order to tightly couple the extension to an event which cannot tolerate the overhead of context switching to an application. By decoupling ASHs they can be run whenever a message arrives, unlike applications. An more venerable example is the monitoring system of Deutsch and Grant [22], which allows code fragments to log various kernel events.

More generally, the decoupled nature of ASHs allow them to be invoked *reactively*, when events happen, rather than when an application executes. This fact forms the basis of a common pattern: reactive incorporation of application knowlege into event handling.

Applications have information that OSes can use to make better decisions. Unfortunately, extracting this information can be difficult. Events must be serviced frequently and quickly. Application context-switching is relatively costly and their runtime difficult to bound at a fine-granularity. However, by decoupling code that can make this decision from the application, it can be tamed and run reactively. Both DPF filters and ASHs fit this pattern. Other examples include the page replacement extensions of Cao et al.' [13], and the messaging systems of Edwards et al. [24] and Fiuczynski and Bershad [32].

One way to view the benefit of this pattern is that decisions profit from more information: the ability to make reactive rather than *a priori* decisions in a dynamic environment like networking can be quite useful. Another way is that that event handling is best done at occurrence, if for no other reason than it removes the overhead of buffering and reduces latency. A final way is that, in general, the more semantics that are known, the faster an operation can be performed, in part because it can be specialized (made more appropriate) based on semantic constraints.

For example, inspired by the observations of Clark and Tennenhouse [18], ASHs can integrate operations such as checksumming, byte swapping, or encryption into the copy from network buffer to application space. The ASH copy engine allows ASHs to provide a code snippet, which it integrates into a specialized copy loop using dynamic code generation. This facility allows ASHs to eliminate duplicate data touching steps.

Unlike packet filters, this ASH ability is entirely motivated by speed, not power. The reason for this difference is that the data copy loop does not involve protection — since all operations are done on the ASH's data — and, therefore, need not be restricted in any way. Packet filters, on the other hand, allow untrusted code to compute protection-critical functions, making restrictions essential.

To summarize: (1) tamed code can be made lightweight and, thus, it can be run in situations where application scheduling is infeasible, and (2) this decoupling allows application semantics to be reactively incorporated into event processing.

6.2.2 Discussion

Source level versus object code level sandboxing. Software fault isolation (SFI) can be done either at the source or at the object code level [59, 89]. ASHs were built using object code SFI, which has the theoretical advantage that it works across languages and compilers, and with pre-compiled code. In retrospect, a source level implementation would have been better. Object-level modification is difficult and extremely non-portable, obviously varying across different architectures and object code formats. Even worse, it also varies across different compiler releases due to the practice of commercial vendors of deliberately changing object code formats in undocumented ways in order to stifle third-party competitors [59].

Source-level SFI, because it is tightly integrated within a compiler, is much simpler to develop. It requires the addition of modest, mostly portable operations done at the level of a compiler's intermediate representation. A secondary benefit of integration is that SFI operations are optimized by the host compiler, unlike object SFI implementations. In contrast, object SFI constantly must fight against the fact that it has lost much lost semantic information. For example, object code modification requires that compiler-generated jump tables be relocated, which can be challenging, since simply finding these tables is difficult, typically requiring compiler-specific heuristics.

An obvious disadvantage of using source-level SFI is that it is specific to one compiler back end and whatever languages its front end(s) consume. In theory, this is important. In practice, operating system software is written in the C programming language. On those rare systems where other languages are used, special support would be required even with object code SFI, since these languages typically use a runtime system, which must be adapted to run in an operating system's restrictive context.

A more serious problem is that a source SFI system typically increases the size of the trusted computing base more than object code SFI system does. A specious counter to this problem is the belief that if the trusted compiler is the same as that used to compile the operating system, then the trusted computing base has not really increased since the correctness of that compiler must already be trusted. However, there is a large difference between compiling a kernel correctly and resisting the malice of clever hackers trying to find a hole in 10-100K lines of compiler, assembler, and dynamic linker code.

The lack of a widely available object code SFI system forced us to "roll our own." Realistically, the reliability of a trusted compiler is most likely better than a object code SFI module we have implemented ourselves.

We no longer use ASHs. In theory, coupling packet arrival to application semantics is profitable. Separate from our work, Edwards et al. [24] provide a way use simple application-provided scripts to direct message placement while Fiuczynski and Bershad [32] provide a fully general messaging system. However, practical technology tradeoffs have made us eliminate ASHs. The three main benefits of ASHs are (1) elimination of kernel crossings, (2) integrated data copying, and (3) fast upcalls to unscheduled processes, thereby reducing processing latency (e.g., of send-response style network messages). On current generation chips, however, the latency of I/O devices is large compared to the overhead of kernel crossings, making the first benefit negligible. The second does not require downloading code, only an upcall mechanism [19]. In practice, it is the latter ability that gives speed. Finally, the presence of DMA hardware makes the data integration ASHs provide irrelevant, since there is no way to add user extensions to the hardware's brute copy.

6.3 XN: efficient disk multiplexing

> Language shapes the way we think, and determines what we can think about. — B. L. Whorf

This section explores language issues in the exokernel's disk multiplexing system, XN, our most "language heavy" subsystem, which contains our most interesting use of downloaded code. We first discuss how to use UDFs to let untrusted file systems track what disk blocks they own, hopefully in enough detail that other practitioners can see how to use UDFs in their domain. We then present a series of insights that have arisen in XN and close with lessons.

6.3.1 Language Evolution

As discussed in Chapter 4, the languages we use to describe client meta data have gone from through four iterations, becoming increasingly general, lower-level, and abstract. We went from an approach that did not use a language at all to one with an expressive declarative description language (which reading file system literature showed as not expressive enough) to our quasi-Turing complete [3] One view of the above evolution is as a struggle to define a universal data layout language. A universal language for most domains requires Turing completeness. Once this fact is realized, it becomes obvious that one needs to provide general-purpose computational primitives. Unfortunately, we only made this connection in hindsight after several years of struggles with disk multiplexing.

6.3.2 Insights

Infinite generality requires Turing completeness. A designer attempting to build an infinitely (or even very general) interface or component set is implicitly striving for Turing completeness. Explicit articulation of this fact makes it clear a solution is to allow clients to customize policies and interface implementations using a Turing complete language rather than, say, a "jumble of procedure flags."

The author belatedly had this insight after struggling with the problem of how to define a completely general set of meta data building blocks and, in fact, only months after coming up with the solution (UDFs) did *why* they solve the problem become clear.

Turing completeness guarantees infinite extensibility. "Solving" extensibility requires showing that any unanticipated use of a system can be implemented. Proving a negative property is hard. A key realization of this chapter is that, when appropriate, Turing completeness provides a way to guarantee that the extensibility problem is solved. For example, the fact that UDFs are (roughly) Turing complete guarantees that that they can describe any computable data layout, anticipated or not.

Transmuting the imperative to the declarative. System implemented functionality is imperative. It determines how to resolve tradeoffs in interface construction: i.e., whether to optimize for latency, throughput, or space. Such predetermination can cause problems when many tradeoffs exist. Downloaded code can be used by a system designer to defer such tradeoffs to clients. By allowing client code to implement functions, the system builder can switch from an imperatively deciding how to implement this function, to declarative testing that client code did so correctly. For example, rather than imperatively deciding how to represent meta data XN declaratively tests that a UDF produces the correct output. This approach has been noticeably easier than the previous struggle to construct a universal data layout language.

The cost of this approach is that testing can be more complex than implementing the functionality (it can also be simpler) and more expensive, though this can be a net win if the algorithm is not on the critical path or grants sufficient power or speed.

Code enables semantic compression. Data representation is important. In a sense, UDFs can be viewed as semantics-exploiting meta data compressors. One could, after all, define a space-inefficient and inflexible but fully general meta data layout. However, UDFs allow representation to be more succinct. For example, much of the meaning of a libFS's meta data is encode in its code, eliminating the need to duplicate this information in the meta data itself (e.g., a libFS "just knows" that certain types of block pointers point to four contigous blocks rather than one). As a more sophisticated example, consider an algebraic relation between blocks such as a file system that allocates blocks at the beginning of every cylinder group. While a predefined data structure would have to list every block, a UDF can encode this knowledge in a function that reads the base block from a instance of meta data and constructs its set:

```
proc owns(meta)
    base = meta  >base_block;
    set = {};
    for i = 0 to number_of_partitions
        set = set U { i * blocks_in_cylinder_group + base };
    return set;
```

[3] While the language used to write UDFs is more-or-less Turing complete, their execution environment is restricted (since UDFs cannot run "too long") as are UDFs (since they must be practical to verify). Similar restrictions will hold for any code downloaded into the kernel, since it must be prevented from at least corrupting kernel data structures. For linguistic we will still refer to such extensions as Turing complete despite this restricted execution environment.

UDFs can make code transparent. A strength of downloaded code is that it can compute its results however it wishes, in ways the underlying system did not anticipate. However, mysterious result computation can also be a liability: users of the code may want to know when it computes a certain output. For example, consider a library file system function, access, that given a principal and inode, indicates whether that principal is allowed to use the inode (access(inode, pid) → bool). Given this function, it is not obvious what values of pid will cause it to return true for a piece of meta data. Thus, an application creating a file controlled using access cannot determine if there exists a special "back door" value of pid that would give others access to its files. UDFs can eliminate this problem. First, we transform this function into one that given an inode produces the set of principles allowed to use it (access(inode) → { set of principles }). (In a sense, we transform access into a function that returns the set of values for which the original access returned true.) Second, we ensure that access is deterministic. Now, at each modification of an inode we can use online testing to ensure that the set of principles associated with it grows or shrinks exactly as it should.

Program verification enables "nestable" extensibility. Because XN verifies the correctness of reference counts, pointers, and meta data interpreters, it allows untrusted implementors to extend an existing file system without compromising its integrity. Thus, it is possible to add an entirely new directory type to a file system and have it point to old types, perform access control on them, etc. without the existing implementation open to malice. We do not know of any other way to achieve this same result.

6.3.3 Lessons

Provide reasonable defaults. UDFs are written in a pseudo-assembly language. A simple virtual machine can be both easy to implement (ours took roughly a day) and small (ours was a few hundred lines of code). The cost, of course, is the unpleasantness of writing assembly-level code. Fortunately, for limited domains, such as packet filters or meta data interpreters, this drawback can be eliminated by hiding such code behind higher-level procedural interfaces, which clients use instead.

Unfortunately, we repeatedly neglected to construct good default libraries after building the base extensible system. As a result, clients typically wrote their code in terms of raw (albeit portable) assembly language, leading to impenetrable code scattered throughout programs. This cycle was self-reinforcing: new programmers that wanted to use the system would look at existing clients, typically not understand what they did, and so cut-and-paste the original code with ad hoc modifications.

The observation that an extensible system must provide good default libraries is neither deep nor unique [11]. Nonetheless it was frequently violated: we did so with packet filters, then with wake up predicates, then with UDFs.

Low-level type systems are useful. From one perspective XN can be viewed as a dynamic type system placed below a file system to catch errors. It served this role well, catching several errors in the C-FFS file system that had escaped the notice of an experienced implementor. In this sense, XN has benefits similar to a low-level typing system such as the Til assembly language developed to catch low-level compiler bugs [85]. One possible use of exokernel technology is to place them below existent operating systems as efficient runtime type checkers.

Fast languages are unnecessary. Writing downloaded code using an efficient language does not hurt. But, at least in the context of an exokernel, it does not seem to help overly much either. UDFs, for example, are written in an interpreted assembly code and wakeup predicates have numerous excess manual address translations. The reason for this is that code used for protection is usually off the critical path, while non-protection code can be placed in the application itself at (perhaps) the cost of an extra system call.

6.4 Protected Methods

This section briefly discusses our most recent use of downloaded code, *protected methods*. Similar to [9], we provides them as an extensible means for applications to implement safe decentralized sharing of state in those cases where the kernel's low-level access control is insufficient. For example, a file system whose directories are mapped to exokernel-protected disk blocks may also require that names within a directory be unique, an invariant inexpressible solely in terms of hardware protection. Using protected methods, directory blocks could be associated with a unique_name method that untrusted library file systems would have to use to allocate names.

This use of downloaded code has little to do with speed. Rather it is intended to solve the problem that distrusting application cannot force one another to obey execution invariants. For example, consider a model where unique_name

is provided in a library, with the admonishment to use it when modifying a directory. Nothing prevents a malicious application from jumping into the middle of the procedure to skip over any invariant checks or, more simply, just writing to the directory block directly.

Mutually mistrusting applications can safely share state by agreeing on code to use, downloading it in the kernel, and accessing the state through this code. Method invocation happens via a system call, forcing execution to begin at a well-defined program counter value. This prevents applications from jumping over guards. Method execution cannot be "hi jacked" by the application manipulating its state via debugging system calls, signals, or page mapping tricks. State exists only in the method's address space. Applications cannot modify it by forging pointers or aliasing virtual addresses. A benefit of this separation is that non-page protection can be readily implemented, as opposed to page granularity memory protection of unrestricted application code. The method cannot write outside its address space. This protects the application from buggy or malicious method code. (Though, the areas where one would trust a method's output, but not trust it to not corrupt the application's state are rare.)

Methods can be used to force the coupling of state modifications, such as forcing the invalidation of a "negative name cache" when a directory entry is allocated. They also help the modification of data structures that span trust boundaries. For example, they can be used to repair file system data structures after a file system crash. Finally, they can provide an easy way to get atomicity.

Protected methods are only one of many ways to provide extensible protection. An alternative is to force all applications to be written in a restricted language and compiled with a trusted compiler. With the advent of languages such as Java, such alternatives may become more palatable.

Another, more traditional way, is to use servers to encapsulate sensitive state. However, because server code is not controlled by the trusted kernel it cannot enforce invariants on it. Thus, server functionality must be trusted completely, and cannot be nested in the same way that methods can be. For example, the kernel can use testing to verify that methods only touches a specific range of bytes in its guarded state, that its modifications preserve pre- and post-conditions, that it is correct, etc.

6.5 Discussion

Data is not a Turing machine. Data is inflexible, but transparent, and its operations (read and write) trivially bounded in cost and guaranteed to terminate. Code is not necessarily any of these three things. The benign characteristics of data can be a relief compared to the uncertainty induced by injecting potentially non-terminating black boxes into a complex operating system. Two specific examples follow.

Information can be communicated by memory (e.g, a flag set when the operating system is allowed to write a disk block) or by code (e.g., a routine called with the disk block asking if it can be written). XN explicitly uses pointers rather than code to track block write orders, despite the lack of flexibility. Code made dependency cycles more difficult to check, and thus, the cost of sharing higher.

It is relatively simple and well understood how to decouple application actions from application execution using the stylized method of buffering. An application that wishes to send a large message on the network can give the buffer for the message to the OS. The OS (or DMA engine) can in turn send it across the network at whatever rate the network supports, irrespective of whether the application that provided the data is currently running. As a result, this pattern of using buffering rather than downloaded code to decouple application actions from scheduling can be seen in all operating systems the author is aware of.

The alternative, having the application explicitly cooperate with the OS via tamed downloaded code, can be accomplished [32, 93] but requires far more machinery than the simple buffer management routines above.

Data has visible transitions. Data, because it is passive, can only be changed by an active entity. Given the right framework (e.g., if applications can only write to data via system calls), this characteristic makes transitions clear and easily coupled to actions or checks. For example, to allow applications to control the order of disk block writes, XN allows them to create dependency chains. Since the pointers used to form chains can only be added using the OS, it is simple to check for cycles. If code was used instead (e.g., a boolean procedure `can_write?`) that was associated with each block), if the code is in any way opaque, such checking becomes more difficult.

Data is passive, control active. From the application's perspective, the passivity of data can be a significant problem: it makes state transitions invisible, requiring polling to track them. Our exokernel provides *wakeup predicates* [4] as a way to conceptually (if not actually) transmute passive data into active events. Wakeup predicates

[4] The exokernel's system for this was conceived, designed and implemented by Thomas Pinckney

are application code snippets downloaded into the kernel, bound to useful memory locations (block I/O flags, timer counters, etc.) and evaluated on various interrupts. When they evaluate to true, the process is awakened.

Code imposes indirection. Placing code between the operating system and its data forces the OS to go through a layer of potentially non-terminating indirection. For example, rather than meta data traversal routines that simply walk down a vector of block pointers, XN requires the use of untrusted iterators and must make provisions to ensure that they do not run too long. Additionally, it cannot simply modify client meta data anymore — since it does not understand their semantics — and, as a result, cannot necessarily do an operation as simple as changing one disk block pointer to another in order to compact the disk.

Further, the indirecting code forces the OS to plan for failure. It must have a contingency plan for when the application code does not update data appropriately or even terminate. Having to rely on a non-trustworthy opponent to do crucial operations can be a practical irritant.

Code hides information. Information can be exchanged from operating system to application using either mapped OS data structures or system calls. The latter interface shields applications from implementation details. However, if the OS does not anticipate the need for a piece of information and encapsulate it within a system call, the client cannot recover it. However, by ripping away this procedural layer layer and exporting data structures (read only) to applications, they can obtain all information, anticipated by the kernel implementor or not. (The potential cost is being tied to a specific implementation.)

Our library operating system's reliance on “wakeup predicates” has driven home the advantages of exposing kernel data structures. Frequently, we have required unusual information about the system. In all cases, this information was already provided by the kernel data structures.

Understanding. A practical problem in using downloaded code is that historically there has been a schism between the compiler and operating system communities. As a result, OS implementors frequently do not understand compilers. They have no equivalent difficulty with data structures.

Alternatives to downloading code for semantics. DPF, ASHs and XN can be viewed as systems to pull application semantics into resource management decisions. However, this is not the only way to get the same effect. The easiest is to upload operating system code into the application. It can then decide how to implement whatever decision it desires. Additionally, it can do so in a Turing complete way, in an unrestricted environment, and with much concern in the operating system about termination and opaqueness issues.

In non-protection situations, the semantics of resources need never be imported into the operating system. The application can instead determine what actions to do, using whatever domain-specific knowledge is important, and then just tell the operating system what to do. This requires constructing interfaces that do declarative checking of an operation rather than imperatively deciding how to do it. For example, consider the problem of writing cached disk blocks to stable storage in a way that guarantees consistency across reboots. Rather than an exokernel deciding on a particular write ordering itself and thus having to struggle with the tradeoffs in scheduling heuristics and caching decisions required to do so well, it can instead allow the application construct schedules, retaining for itself the much simplified task of merely checking that any application schedule gives appropriate consistency guarantees. Application of this methodology enables an exokernel to leave library operating systems to decide on tradeoffs themselves rather than forcing a particular set, a crucial shift of labor.

One of the games we have played frequently in an exokernel is determining how to construct interfaces where an application “just knowing” what is appropriate can be expressed as a kernel action.

6.6 Related Work

There have been a number of papers that evaluate different downloading code mechanisms. Small and Seltzer compare several extension techniques [80], Bershad et al. [9] and Pardyak and Bershad describe [71] different aspects of the SPIN operating system's extensibility framework. This chapter complements this prior work. Stallman [81] and Borenstein and Gosling [11] discuss language issues in the context of the EMACS text editor. The discussion is largely complementary, since extensions in this context are trusted, but there are overlapping lessons (the most painful to rediscover that reasonable defaults must be provided).

Both SPIN [9, 32, 71] and Vino [79, 80] are two other extensible operating systems that use downloaded code.

Code motion has been a recurrent theme in operating systems since inception [41, 22]. Micro-kernels are an attempt to move operating system code out of the harsh environment of the kernel into the more genteel context of processes [41, 53, 2, 77, 84]. Virtual machines [20] similarly move operating system code to application-level.

Interestingly, this is the single feature they change about the OS: all hardware details are emulated faithfully. Most modern operating systems provide ways to dynamically download load device drivers.

Several hints for when to download code can be found in Lampson [52]. A useful insight is that downloading is simply an example of higher-order function programing (or in systems languages, the use of function pointers rather than flags as parameters). The main difference is that code is being shipped across trust boundaries rather than, for instance, library interfaces. Thus, it appears possible to take some of the ideas from these mature areas "whole cloth." Sussman and Abelson [1] is a classic text.

The parallel community has long considered the idea of "function shipping" for speed — e.g., to bring computation closer to data, to be able to migrate computations for load-balancing, etc. Some of the insights from this use can be applied to operating systems. Similarly, the distributed systems community has shipped code as well. Java applets are a topical example [40]. Tennenhouse and Weatherall have proposed to use mobile code to build Active Networks [86]; in an active network, protocols are replaced by programs, which are safely executed in the operating system on message arrival. Curiously, in contrast to our experience, most uses of mobile code in an Active Network seem to be to improve efficiency, rather increase power.

Chapter 7

Conclusion

> But in our enthusiasm, we could not resist a radical overhaul of the system, in which all of its major weaknesses have been exposed, analyzed, and replaced with new weaknesses. — Bruce Leverett, "Register Allocation in Optimizing Compilers"

This chapter discusses possible ways that an exokernel approach could fail, lessons learned in building our exokernel systems, and conclusions.

7.1 Possible Failures of the Architecture

> Doubt 'til thou canst doubt no more...doubt is thought and thought is life. Systems which end doubt are devices for drugging thought. —- Albert Guerard

In our mind, the remaining serious questions about the exokernel architecture are sociological ones rather than technical. We list five possible failures of the architecture once it moves from our coddling laboratory into the "real world:"

1. Application writers do not deal well with the freedom they have and become tied too closely to a particular exokernel implementation, preventing upgrades and slowing, rather than improving, system evolution. While adherence to standard interfaces and good programming practices *should* prevent this type of failure (given that these techniques have a venerable track record of doing so in other domains), it remains to be demonstrated if they suffice for exokernels.

2. Commoditization of operating system software makes operating system research irrelevant. Commoditization has already severely restricted the viability of new OS interfaces. The dominance of a few OSes, and the increasing cost of implementing them, may restrict the viability of new implementations of these interfaces as well. If so, then much of the innovation potential of an exokernel will be lost.

3. The technical ability to innovate does not lead to any more innovation than on traditional systems. The exokernel architecture is based on a partially-sociological assumption: that making OS innovation easier and less costly will lead to a vast improvement in innovation. This may well be a false assumption. For instance, innovation may already be "easy enough" for those who care to do it.

4. An exokernel, by migrating most OS code to libraries, removes the "single point of upgrade" characteristic of current systems. This can be an advantage, since applications do not have to wait for the central OS to upgrade but can instead do so themselves. However, it can also impede progress by making improvements harder to disseminate.

5. Users will not switch operating systems. An OS forms the primal mud on which systems are built. Changes to it have far reaching impact. Computer system users have thus demonstrated an understandable reluctance to alter it. It may be that the advantages of an exokernel system do not proof sufficient to lead to such a switch.

> Fortunately, an exokernel does not require "whole cloth" adoption for success. It appears that many exokernel interfaces, especially those related to I/O, can be grafted on to existing systems, with little loss in performance. Normal applications would the existing OS interfaces as a default, while more aggressive applications would have the power to control important decisions. [1]

While we have confidence that the exokernel's technical advantages will allow it to transcend these potential pitfalls, it must still demonstrate that it does.

7.2 Experience

Over the past three years, we have built three exokernel systems. We distill our experience by discussing the clear advantages, the costs, and lessons learned from building exokernel systems.

7.2.1 Clear advantages

Exposing kernel data structures. Allowing libOSes to map kernel and hardware data structures into their address spaces is a powerful extensibility mechanism. (Of course, these structures must not contain sensitive information to which the application lacks privileges.) The benefits of mapping data structures are two-fold. First, exposed data structures can be accessed without system call overhead. More importantly, however, mapping the data structures directly allows libOSes to make use of information the exokernel did not anticipate exporting.

Because exposed data structures do not constitute a well-defined API, software that directly relies on them (e.g., the hardware abstraction layer in a libOS) may need to be recompiled or modified if the kernel changes. This can be seen as a disadvantage. On the other hand, code affected by changes in exposed data structures will typically reside in dynamically-linked libOSes, so that applications need not concern themselves with these changes. Moreover, most improvements that would require kernel modification on a traditional operating systems need only effect libOSes on exokernels. This is one of the main advantages of the exokernel, as libOSes can be modified and debugged considerably more easily than kernels. Finally, we expect most changes to the exokernel proper to be along the lines of new device drivers or hardware-oriented functionality, which expose new structures rather than modify existing ones.

In the end, some aggressive applications may not work across all versions of the exokernel, even if they are dynamically linked. This problem is nothing new, however. A number of UNIX programs such as top, gated, lsof, and netstat already make use of private kernel data structures through the kernel memory device `/dev/kmem`. Administrators have simply learned to reinstall these programs whenever major kernel data structures change.

The use of "wakeup predicates" has forcefully driven home the advantages of exposing kernel data structures. Frequently, we have required unusual information about the system. In all cases, this information was already provided by the kernel data structures.

The CPU interface. The combination of time slices, initiation/termination upcalls, and directed yields has proven its value repeatedly. (Subsequent to our work, others have found these primitives useful [35].) We have used the primitives for inter-process communication optimization (e.g., two applications communicating through a shared message queue can yield to each other), global gang-scheduling, and robust critical sections (see below).

Libraries are simpler than kernels. The "edit, compile, debug" cycle of applications is considerably faster than the "edit, compile, reboot, debug" cycle of kernels. A practical benefit of placing OS functionality in libraries is that the "reboot" is replaced by "relink." Accumulated over many iterations, this replacement reduces development time substantially. Additionally, the fact that the library is isolated from the rest of the system allows easy debugging of basic abstractions. Untrusted user-level servers in microkernel-based systems also have this benefit.

7.2.2 Costs

Exokernels are not a panacea. This section lists some of the costs we have encountered.

Exokernel interface design is not simple. The goal of an exokernel system is for privileged software to export interfaces that let unprivileged applications manage their own resources. At the same time, these interfaces must offer rich enough protection that libOSes can assure themselves of invariants on high-level abstractions. It generally takes several iterations to obtain a satisfactory interface, as the designer struggles to increase power and remove unnecessary

[1] John Jannotti in our group has begun the process of integrating exokernel disk and network interfaces into Linux.

functionality while still providing the necessary level of protection. Most of our major exokernel interfaces have gone through multiple designs over several years.

Information loss. Valuable information can be lost by implementing OS abstractions at application level. For instance, if virtual memory and the file system are completely at application level, the exokernel may be unable to distinguish pages used to cache disk blocks and pages used for virtual memory. Glaze, the Fugu exokernel, has the additional complication that it cannot distinguish such uses from the physical pages used for buffering messages [60]. Frequently-used information can often be derived with little effort. For example, if page tables are managed by the application, the exokernel can approximate LRU page ordering by tracking the insertion of translations into the TLB. However, at the very least, this inference requires thought.

Self-paging libOSes. Self-paging is difficult (only a few commercial operating systems page their kernel). Self-paging libOSes are even more difficult because paging can be caused by external entities (e.g., the kernel touching a paged-out buffer that a libOS provided). Careful planning is necessary to ensure that libOSes can quickly select and return a page to the exokernel, and that there is a facility to swap in processes without knowledge of their internals (otherwise virtual memory customization will be infeasible).

7.2.3 Lessons

Provide space for application data in kernel structures. LibOSes are often easier to develop if they can store shared state in kernel data structures. In particular, this ability can simplify the task of locating shared state and often avoids awkward (and complex) replication of indexing structures at the application level. For example, Xok lets libOSes use the software-only bits of page tables, greatly simplifying the implementation of copy on write.

Fast applications do not require good microbenchmark performance. The main benefit of an exokernel is not that it makes primitive operations efficient, but that it gives applications control over expensive operations such as I/O. It is this control that gives order of magnitude performance improvements to applications, not fast system calls. We heavily tuned Aegis to achieve excellent microbenchmark performance. Xok, on the other hand, is completely untuned. Nevertheless, applications perform well.

Inexpensive critical sections are useful for LibOSes. In traditional OSes, inexpensive critical sections can be implemented by disabling interrupts [10]. ExOS implements such critical sections by disabling software interrupts (e.g., time slice termination upcalls). Using critical sections instead of locks removes the need to communicate to manage a lock, to trust software to acquire and release locks correctly, and to use complex algorithms to reclaim a lock when a process dies while still holding it. This approach has proven to be similarly useful on the Fugu multiprocessor; it is the basis of Fugu's fast message passing.

User-level page tables are complex. If page tables are migrated to user level (as on Aegis), a concerted effort must be made to ensure that the user's TLB refill handler can run in unusual situations. The reason is not performance, but that the naming context provided by virtual memory mappings is a requirement for most useful operations. For example, in the case of downloaded code run in an interrupt handler, if the kernel is not willing to allow application code to service TLB misses then there are many situations where the code will be unable to make progress. User-level page tables made the implementation of libOSes tricky on Aegis; since the x86 has hardware page tables, this issue disappeared on Xok/ExOS.

7.3 Conclusion

Our inventions mirror our secret wishes. Lawrence Durrell (1912-1990) “Mountolive” (1959)

This thesis proposes and evaluates the exokernel operating system architecture. An exokernel gives untrusted application code as much safe control over resources as possible. It does so by separating management from protection. All functionality necessary for protection resides in the exokernel, control over all other aspects is given to applications. Ideally, applications can safely and efficiently perform any operation that a privileged operating system can. Thus, unlike traditional systems, on an exokernel system, OS software becomes: (1) unprivileged, (2) able to co-exist with other implementations, (3) modifiable and deployable by orders of magnitude more programmers. We hope that this organization significantly improves operating system innovation.

This thesis has discussed both the exokernel architecture, and how to apply its principles in practice, drawing upon examples from two exokernel systems. These systems give significant performance advantages to aggressively-specialized applications while maintaining competitive performance on unmodified UNIX applications, even under

heavily multi-tasked workloads. Exokernels also simplify the job of operating system development by allowing one library operating system to be developed and debugged from another one running on the same machine. The advantages of rapid operating system development extend beyond specialized niche applications. Thus, while some questions about the full implications of the exokernel architecture remain to be answered, it is a viable approach that offers many advantages over conventional systems.

Chapter 8

XN's Interface

This Chapter describes the public and privileged XN system call interface.

A number of routines expect a device number (an integer of type dev_t), which names an active XN-controlled disk. This name is implicit in the disk address type, da_t, a 64-bit integer that encodes the device name, disk block, and byte offset within the disk block. A device corresponds to some range of disk blocks, a freemap that tracks these blocks, and a "root catalogue," which is a persistent table that libFSes install types and file system roots into.

In general, any XN system call fails if: (1) a libFS-supplied capability is insufficient, (2) a libFS-supplied pointer is bogus (i.e., not readable or, for modifications, not writeable), or (3) a libFS-supplied name is bogus (e.g., an invalid disk block name, an invalid root catalogue entry character string, etc.). To save space, we do not mention these errors further.

We elide the details of mapping XN data structures and buffer cache entries, since they are specific to the hosting OS's virtual memory interface rather than XN itself.

A Privileged system calls

Only the kernel or privileged applications can use the following system calls.

A.1 XN initialization and shutdown

The following three routines are used to initialize and cleanly shutdown an XN-controlled disk.

xn_err_t sys_xn_init(dev_t dev);

> Initialize XNdisk dev. Fails if the disk was not cleanly shutdown, in which case the XN disk reconstruction program must be run in order to find all XN invariant violations.

xn_err_t sys_xn_shutdown(dev_t dev);

> Clean shutdown of an disk dev. Fails if any blocks in the buffer cache contain violations. Before it can proceed the invoker must iteratively flush the buffer cache to remove these blocks.

dev_t sys_xn_format(void);

> Prepare a new disk for XN; returns the device number of the disk. Currently, it always succeeds.

A.2 Reconstruction

The following two functions are used by privileged reconstruction programs:

db_t sys_db_alloc(dev_t dev, db_t db, size_t n);

> Forces the extent [db, db+n) on disk dev to be taken off of the freelist.

xn_err_t sys_xn_mod_refcnt(da_t da, int delta);

> Alters da's reference count by delta.

Currently, the error log routines are in flux. We elide them here.

B Public system calls

The following system calls are intended for use by any libFS.

B.1 Creating types

The following three system calls are used to create the types for a new new libFS.

xn_err_t sys_install_mount(dev_t dev, char *name, db_t *db, size_t nelem, xn_elem_t t, cap_t c);

> Allocate the extent [db, db+nelem*sizeof t) on disk dev and associate it with name in the root catalogue. If db's value is 0, then the kernel decides what extent to allocate and writes the allocated block into db. LibFSes use this system call to install both install file system roots and types. On reboot, the XN reconstructor loads the catalogue and traverses file systems from these roots, garbage collecting and performing consistency checks. The entry can be modified for applications that possess cap. The call fails if any block in the extent has already been allocated, if type is invalid, name or db are not readable.

xn_err_t sys_type_import(dev_t dev, char *type_name);

> Converts the root catalogue entry name in dev's root catalogue from a disk block extent to an actual type. It fails if name does not exist, the blocks are not "raw" disk blocks (i.e., of type XN_DB), or the contents of the extent form an invalid type.

xn_err_t sys_reserve_type(dev_t dev, char *name);

> Reserve a slot for an as-yet-unspecified type in dev's root catalogue. This call is used to construct mutually recursive types.

To create a new type, the libFS performs the following four steps:

1. Allocates space for it on disk and installs it in the root catalogue using sys_install_mount. The value of type in this installation is XN_DB, indicating the blocks are simple disk blocks. If the type is involved part of a mutually recursive specification (which can happen when the type is a composite type), the libFS can use sys_reserve_type to reserve the type names for these other types, as well as getting their type id (an integer assigned by the kernel), which is needed by the type's owns UDF to compute the typed block set that it outputs.
2. Initializes these blocks using sys_xn_writeb (discussed below) to the contents in ths types "type structure," which holds the owns UDF, the refcnt UDF, and other information.
3. Writes these blocks back to disk (the following step fails if they are dirty).
4. Finally, it uses sys_type_import to convert the blocks from generic blocks to an XN recognized type. sys_type_import performs consistency checks on the type data structure (e.g., that the contained UDFs are deterministic, that the partitions are sensible) and, if the checks succeed, changes the type cataglogue entry to be of type XN_TYPE rather than XN_DB.

At this point, the libFS can create and point to blocks of the new type.

B.2 Creating and deleting file system trees

To create a new file system, the libFS installs the root of the tree using sys_install_mount. To load an already existing one from disk into the buffer cache it can use the following two functions:

xn_err_t sys_xn_mount(dev_t dev, struct root_entry *r, char *name, cap_t c);

> Bootstrap the file system tree by inserting the type name's block extent into the buffer cache: getting the types for the rest of the file system tree can be done by recursively applying the associated owns UDF to the root and its children. The call writes the root catalogue entry into r, and then annotates the associated entries in the buffer cache with the given type. Fails if name name does not exist, or the extent pointed to by the root catalogue entry is not in the buffer cache registry.

```
/* specify xn operations that involve UDFs. */
struct xn_op {
        /* Specify what extent will be allocated/freed/read. */
        struct xn_update {
                size_t own_id;          /* id of the udf to run. */
                cap_t   cap;            /* capability to use for access. */

                db_t db;          /* base    all objects are sector aligned. */
                size_t nelem;     /* number of elements of type. */
                xn_elem_t type; /* type */
        } u;

        /*
         * Specify update to a piece of meta  data.  Semantics:
         *       memcpy((char *)meta + offset, addr, nbytes);
         * Ignored for reads.
         */
        struct xn_m_vec {
                size_t   offset;  /* what offset in the type */
                void    *addr;  /* ptr to value to copy there */
                size_t   nbytes;
        } *mv;
        size_t n;               /* number of elements in mv */
};
```

Figure 8-1: Metadata operation structure.

xn_err_t sys_type_mount(dev_t dev, char *type_name);

Similar to sys_xn_mount, except that it annotates a a type_name's cached blocks as holding an XN type. It fails if the given name does not exist in the root catalogue, is not a type, or if the blocks are not in the buffer cache.

To use a file system tree, the file system root and any types it needs must first be loaded into the buffer cache using sys_xn_read_and_insert (discussed below). Then, sys_xn_mount annotates these disk blocks with the type given by their root catalogue entry (for the root, its synthetic libFS type, for types themselves XN_TYPE).

Types and roots can be removed from the root catalogue using:

xn_err_t sys_uninstall_mount(dev_t dev, char *name, cap_t c);

Delete root catalogue entry name on dev. Fails if name is a file system root that contains child pointers, or is a type with cached blocks. If a type is deleted that is used by some non-cached meta data instance, then that meta data's owns function will fail (as will the system call using owns). A better solution would be to count how many blocks of a given type exist and only allow the type to be deleted when there are no more blocks of its type.

B.3 Buffer cache operations

The structure xn_op is used for many of the remaining system calls, usually to specify extents, owns partition id's, and (for modification) the byte range(s) to modify. Figure 8-1 gives its ANSI C representation.

The following two routines are used to bring blocks into the buffer cache and prepare them for use and delete them:

xn_err_t sys_xn_read_and_insert(dev_t dev, db_t db, size_t nblocks, xn_cnt_t *cnt);

```
struct xn_iov {
        xn_cnt_t *cnt;              /* Decremented on every successul io*/
        struct xn_ioe {
                db_t db;
                size_t nblocks;
                void *addr;         /* mem to write from/to. */
        } iov[1];
        size_t n_io;               /* number of entries. */
};
```

Figure 8-2: I/O vector structure.

Read [db, db+nblocks) from disk dev into the buffer cache. cnt is incremented each time a block is brought in. Fails if the extent does not fit in the buffer cache.

xn_err_t sys_xn_insert_attr(da_t parent, struct xn_op *op);

Install the type of the buffer cache entry using the parent block named by parent. Allocation, deletion and reading and writing all perform this call implicitly since they require registry entries. Fails if the extent specified in op is not guarded by the sub-partition specified in op.

xn_err_t sys_xn_delete_attr(dev_t dev, db_t db, size_t nblocks, cap_t cap);

Remove the buffer cache entries for [db, db+nblocks) associated with disk dev. Fails if removing the entry would lead to an invariant violation or if the entry is in use by some other application.

Importantly, libFSes can fetch any block extent into the buffer cache. They only need to locate the blocks parent (and thus its type and access control mechanism) before actually reading or writing the cached blocks.

The following two functions write buffer cache entries back to disk:

xn_err_t sys_xn_writeback(dev_t dev, db_t db, size_t nblocks, xn_cnt_t *cnt);

Writes [db, db + nblocks) back to disk dev. Each write decrements cnt.

xn_err_t sys_xn_writebackv(dev_t dev, struct xn_iov *iov);

Writes a list of these extents back to disk dev. On some hardware disks, this "gather" mechanism enables more sophisticated scheduling. Figure 8-2 gives the ANSI C representation for xn_iov. Note that neither of these two system calls require access checks: They only fail if the extent is invalid or contains blocks tainted with some violation.

The following two routines read and write the bytes in buffer cache entries. No byte range read by a UDF can be modified using these routines.

xn_err_t sys_xn_writeb(da_t da, void * src, size_t nbytes, cap_t cap);

Writes [src, src + nbytes) into the cached blocks for [da, da + nbytes). All of the bytes must be in memory. Fils if any block of the extent is not in core.

xn_err_t sys_xn_readb(void * dst, da_t da, size_t nbytes, cap_t cap);

Reads [da, da + nbytes) into [dst, dst + nbytes). Fails for reasons identical to above, with the additional restriction that [dst, dst + nbytes) must be writable.

B.4 Metadata operations

The following routines are used to allocate blocks, form edges to existing blocks, delete edges to existing blocks, and change a block's type:

xn_err_t sys_xn_alloc(da_t da, struct xn_op *op, unsigned alloc);

Allocates the extent [db, db + (sizeof type * nelem) given by op and writes the modifications contained in op into the parent metadata, da. alloc indicates whether the block should be zero filled. The call fails if: (1) the extent cannot be allocated; (2) da is not in core; (3) op contains bogus data (its partition id is invalid or the modification vector has bogus byte ranges); or (4) the UDF returns the wrong data.

xn_err_t sys_xn_free(da_t da, struct xn_op *op);

Frees the extent [db, db + (sizeof type * nelem) given by op. If the extent's has a reference count greater than one, then the reference count is decremented. Otherwise it is placed on the freelist. Fails for similar reasons to sys_xn_alloc.

xn_err_t sys_xn_add_edge(da_t src, struct xn_op *src_op, da_t dst);

Forms an edge from src to dst and increments dst's reference count. Fails if: (1) neither node is in core (src must be modified to add a pointer, dst modified to increment its refernce count); (2) src's owns UDF fails, or dst's refcnt UDF fails; or (3) src_op specifies bogus modifications.

xn_err_t sys_xn_set_type(da_t da, int ty, void *src, size_t nbytes, cap_t cap);

Change the type of a type union instance: XN metadata types can be "unions," which means they can be dynamically converted among a series of listed types. Currently, the type field must be "nil" (i.e., the type has just been initialized). Extending the system call to convert between existant types would not be difficult (it only requires retesting that the new type's owns function emits the same blocks as the old one).

B.5 Reading XN data structures

The following six functions allow applications to read various XN bookkeeping data structures. All fail if the given memory is not writable. If the used data structures are mapped into the libFS's addresses space, these calls are simple library calls.

xn_err_t sys_xn_read_attr(dev_t dev, void * dst, db_t db, cap_t cap);

Reads the registry attribute for db on disk dev into dst. Fails if db is not in core.

xn_err_t sys_xn_read_catalogue(dev_t dev, struct root_catalogue *c);

Read dev's root catalogue into c.

xn_err_t sys_dev_list(dev_t *devl, int ndevs);

Stores the enumerated list of active devices into devl, which is ndevs big. Fails if the destination is not large enough.

xn_err_t sys_xn_info(dev_t dev, db_t* r_db, db_t* f_db, size_t* f_nbytes);

Reads the block address that holds the root catalogue into r_db, the block address of the freemap into f_db and the size of the freemap into nbytes.

db_t sys_xn_findfree(dev_t dev, db_t hint, size_t nblocks) ;

Returns the first free extent (starting at hint, hint+nblocks)) found on dev. If hint is 0, then XN makes its own decision about where to start searching.

xn_err_t sys_xn_list_writeable(dev_t dev, db_t *dbv, size_t *n);

Reads a the list of dirty but writable entries for device dev in the buffer cache into dbv (whose size is given by n). Typically this is used by two programs. A "syncer" deamon that flushes back dirty blocks, and by any shutdown application that needs to flush all blocks back to disk. The latter locks the buffer cache and iteratively flushes blocks until the empty list is returned.

db_t sys_root(dev_t dev);

Return the block address of dev's XN-created "superblock" (i.e., the block that holds the pointers to XN's per-device bookkeeping data structures).

B.6 File system-independent navigation calls

The following routines allow program to navigate any libFS meta data. They are used, for example, by our disk reconstruction program.

xn_err_t sys_xn_get_refcnt(da_t da, cap_t cap);

Returns da's reference count.

xn_err_t sys_xn_get_nowns(da_t da, cap_t cap);

Returns the number of owns partitions in da's type.

xn_err_t sys_xn_udf_enum(da_t da, size_t owns_b, size_t owns_e, struct xn_update *ups, size_t n, cap_t cap);

Computes a list of of all blocks controlled by da's partitions owns_b through owns_e. These are written into the vector ups (whose size is given by n). To ensure that the list does not get "too big" and the system call does not run "too long" (both of which are OS dependent) the enumeration may be broken up by XN into multiple passes.

Chapter 9

Aegis' Interface

In general, any system call fails if: (1) a libOS lacks permissions for the operation, (2) a libOS-supplied pointer is bogus (i.e., not readable or, for modifications, not writeable), or (3) a libOS-supplied name is bogus (e.g., an invalid page name, packet filter id, etc.). To save space, we do not mention these errors further.

.7 CPU interface

Figure .7 presents the ANSI C representation of an Aegis time slice. The routines to allocate, yield, and free time slices are given below.

```
int ae_s_alloc(int pid, int n);
```

Allocate time slice `n` and give environment `pid` access to it. Fails if `n` is not free, or the current process lacks write access to `pid`.

```
int ae_s_free(int slice);
```

Free time slice `slice`. Fails if the current process lacks write access to the owning environment.

```
int ae_yield(int pid);
```

Yield the remainder of the current time quantum to `pid`.

```
int ae_donate(int pid);
```

Permenantly donate current time slice to process `pid`.

.8 Environments

Figure .8 presents the ANSI C representation of an Aegis environment, which is used to store the hardware information needed to execute a thread of control in a virtual address space, along with resource accounting information.

```
int ae_e_alloc(int n, struct env *e);
```

Allocate environment `n`. Application fills in the guarenteed mappings, exception handlers in the environment structure (given in `e`). Aegis checks that any given physical addresses are allowed an sensible.

```
int ae_e_ref(int pid);
```

Add a refrence from the current process to environment `pid`. Fails if the current process lacks permissions.

```
int ae_e_unref(int pid);
```

Unreference environment `pid`. If there are no other outstanding references, the env is freed. Fails if the current process does not have read permission to the environment.

```
/* Time  slice;  controls an ordered scheduling quanta.  It is donated
 * permenentely by synchronous IPC.  It can be donated temporarily by
 * yielding (via ae_yield) or by asynchronous IPC's (via ae_asynch_ipc).
 * Ownership of the time  slice can be changed by any process that has
 * the owning environments capability.
 *
 * Time  slices are initiated via an upcall to application space and
 * revoked in the same manner; this interaction allows applications
 * to control important context  switching operations (e.g., this
 * functionality is sufficient to implement scheduler activations).
 * The price of this functionality is that an application must be
 * prevented from ignoring revocation interrupts.  The current method
 * is to record the number of timer interrupts a process has ignored
 * in 'ticks' and when this value exceeds a predefined threshold
 * killing the process.  In a more mature implementation we would simply
 * context switch the application by hand.
 */
struct slice {
        char pad[12];
        struct env         *e;     /* associated environment (null if no one) */
        unsigned short  next,prev;      /* forward and back pointers */
        int ticks;          /* ticks consumed in interrupts */
        int epc;
};
```

Figure 9-1: Aegis time-slice representation

```
/*
 * Environment: this is the most complicated entity in aegis.  An
 * environment is basically a process: it defines the program counters
 * to vector events to and serves as a resource accounting point.
 */
struct env {
        /* pointers to exception "save areas" where Aegis stores
         * active registers. */
        addr_t          tlbx_save_area,
                        genx_save_area,
                        intx_save_area;

        /* exception handlers. */
        handler_t       xh[NEX], /* sync exception handlers */
                        epi,    /* epilogue code */
                        pro,    /* prologue code */
                        init;   /* initial code that is jumped to */
        /*
         * The following four fields are clustered to be on
         * the same cache line.
         */
        signed char     cid;    /* address space identifier */
        unsigned char   envn;   /* environment number */
        short           tag;    /* 11 bit tag */
        handler_t       gate;   /* ipc entry point */
        unsigned        status; /* status register */
        struct tlb      xl[MAXXL];      /* guarenteed translations */
        struct ae_intr_queue *iq;       /* interrupt queue */
};
```

Figure 9-2: Environment structure

int ae_e_add(int pid, int n);

Give pid access to environment n. Only gives write access at the moment. Fails if the current process does not have access to pid.

int ae_e_free(int pid);

Free all resources consumed by environment pid.

int ae_e_write(int n, size_t offset, void *p, size_t nbytes);

Modify the byte range [offset, offset+nbytes) in the environment n's application-specific data region.

.9 Physical memory

The following three routines allocate, deallocate, and share physical pages.

int ae_p_alloc(int n);

Allocate page n. Fails if page is already allocated.

int ae_p_unref(int n);

Remove reference to page n. If no other process has a reference to the page it is deallocated. Fails if the current process lacks permissions.

int ae_p_add(int prot, int pid, int n);

Give process pid access to page n with protections prot. Fails if the current process lacks appropriate permissions.

.10 Interrupts

To make interrupt handling efficient, Aegis places interrupt notifications in a user-space interrupt queue that the libOS can read and modify directly. The structures used for this are given in Figure 9-3.

int ae_read_exposed_info (struct exposed_info *i);

Read out offsets and lengths of exposed kernel data structures in preperation for mapping their associated pages. Figure 9-4 presents the data structures that can be mapped.

int ae_getrate(void);

Get the system's clock rate.

int ae_gettick(void);

Get the current clock tick.

int ae_cacheflush (void *ptr, int sz, int flags);

Flush the address range [ptr, ptr+sz) out of the cache.

int ae_memcpy(int dst_pfn, int src_pfn, int sz, int cache);

Copy pages src_pfn, src_pfn+sz) to the contigious page range starting at dst_pfn, bypassing the TLB. cache indicates whether the copy should also bypass the cache.

int ae_memset(int dst_pfn, int offset, char b, int sz);

Set memory [dst_pfn+offset, dst_pfn+offset+sz) to b. Bypasses the TLB.

int ae_fpu(int flag);

Enable/disable floating-point unit.

int ae_pid(void);

Get the current pid.

```
struct ae_interrupt {
        handler_t  h;           /* handler to jump too. */

        /* Up to three int  specific arguments.  Used to parameterize interrupts.*/
        void *arg1;
        void *arg2;
        void *arg3;

        /* Saved values to give application scratch registers
           epc is set to zero if no return point.  */
        unsigned epc;   /* holds value that we were interrupted at. */
        unsigned a0;    /* holds a0 register */
        unsigned a1;    /* holds a1 register */
        unsigned a2;    /* holds a1 register */
};

/*
 * circular interrupt queue: is provided at user  level so that applications
 * can control interrupt handling efficiently.
 *
 * Useful facts:
 *      1. pending == 0    > q is empty
 *      2. tail points to the first entry to dequeue.
 *      3. full: pending == sz
 *      4. head    points to first empty entry.
 *
 * To consume an interrupt:
 *      1. increment tail % AE_INTQ_SZ
 *      2. decrement pending.
 *      3. renable the interrupt type.
 */
struct ae_intr_queue {
        /*
         * This flag counts the number of interrupts pending while the
         * process is running with interrupts disabled.
         */
        unsigned short pending;

        /*
         * Number of overflow ints for each type.  (simple way to allow
         * resource  specific recovery.)
         */
        unsigned short overflow[AE_NINTS];
        unsigned short overflow_p;      /* set to 1 if there is overflow. */

        unsigned        mask;           /* 0   > disabled, 1   > enabled. */

       /*
         * Handlers, two for each interrupt.  We make application
         * explicitly check whether it was running or not.  Could
         * have two handlers, where  handler 0 is set when process
         * was not executing, handler 1 is set when it was.
         */
        handler_t       h[AE_NINTS];

        /*
         * Circular queue; when it runs out, we write an overflow interrupt.
         */
        unsigned char head;
        unsigned char tail;
```

```
struct exposed_info {
    unsigned int start_page;   /* which phys page the info starts on */
    int num;                   /* how many pages */

    /* timing related structures */

    char *clock_tick_start;
    int clock_tick_size;
    char *clock_rate_start;
    int clock_rate_size;

    /* page info */
    char *p_refcnt_start, *p_acl_start, *p_map_start;
    int p_refcnt_size, p_acl_size, p_map_size;

    /* env info */
    char *e_refcnt_start, *e_acl_start, *e_map_start, *env_start;
    int e_refcnt_size, e_acl_size, e_map_size, env_size;

    /* time slice info */
    char *s_refcnt_start, *s_acl_start, *slice_start, *s_map_start;
    int s_refcnt_size, s_acl_size, slice_size, s_map_size;

    /* context ids info */
    char *cmap_start, *c_map_start;
    int cmap_size, c_map_size;

    /* stlb */
    char *stlb_start;
    int stlb_size;
};
```

Figure 9-4: Structure used to hold where each exposed kernel data structure begins and its size. By default, each environment has read access to the pages containing these structures.

```
/* Structure to hold recv messages. */
struct ae_recv {
        int n;  /* number of entries */        struct rec {
                int sz;
                void *data;
        } r[MAXPKTS];
};
```

Figure 9-5: Network packet receive structure

.11 Networking

Aegis provides calls to insert and delete packet filters. It also provides support for Ethernet and AN2 OTTO chips. For simplicity, we only provide the former's interface.

The following two functions install and delete packet filters. Filters can be bound to an ASH (libOS messaging code downloaded into the kernel), which is run when the filter matches. Otherwise the packet will be handled by the libOS.

`int ae_dpf_install(void *filter, int sz, int ash_id);`

Installs filter (of sz bytes) and binds it to the ASH ash_id. Fails if the size is too large, or the filter overlaps with another filter and the current process has insufficient permissions to override it.

`int ae_dpf_delete(int id);`

Delete filter id. Fails if the process lacks permission to do so.

The following four functions are used to send and receive messages.

`int ae_eth_poll(int fid, struct ae_recv *recv, volatile int *addr);`

Install a receive structure, recv, to handle packets arriving for filter fid; receive notification via polling. addr points to a counter that is incremented by the received message size. Figure 9-5 shows the ANSI C representation of the receive structure. Fails if there are already too many enqueued receive structures.

`int ae_eth_send(void *msg, int sz);`

Send msg (of sz bytes) on Ethernet.

`int ae_eth_sendv(struct ae_recv *recv);`

"Gather" send of the message specified by recv: Aegis copies the message into a contigous outgoing buffer (the hardware we run on lacks support for DMA).

`int ae_eth_info(addr_t addr);`

Get Ethernet address.

.12 TLB manipulation

Aegis uses a software TLB [7] to increase the effective hardware TLB size. The STLB has 8192 entries and is a direct mapped hash table; it also has an 8-entry fully-associative overflow buffer (similar to a "victim cache" [47]). Figure 9-6 gives the ANSI C representation of an STLB entry. Figure 9-7 gives the MIPS assembly code to load an entry from the STLB into the hardware TLB within the TLB handler. LibOSes can map the STLB read-only into their address spaces. LibOS modifications of the hardware TLB are propogated to the STLB. The libOS program counter that Aegis vectors TLB misses to is given in the libOSes environment structure.

The following routings, read, insert and modify TLB entries. `int ae_tlbwr(addr_t va, struct lo lo);`

```
struct stlb {
    /* STLB tag: the contents of the TLB context register (c0_tlbctx) */
    unsigned    :2,
                vpn:19, /* bad virtual page number */
                tag:11; /* 11 bit tag associated (pseudo  randomly) with each process */
    /* TLB entry */
    unsigned    :8,     /* reserved */
                g:1,    /* Global: TLB ignores the PID match req */
                v:1,    /* Valid: if not set, TLBL or TLBS miss occurs*/
                d:1,    /* Dirty */
                n:1,    /* Non  cacheable. */,
                pfn:20; /* Page frame number */
};
```

Figure 9-6: STLB structure

Insert TLB entry for virtual address va. lo holds the physical page number, protection information, and various software tags. Figure 9-8 gives its ANSI C representation.

void ae_tlbrprotn(addr_t va, int len);

Read-protect the region [va, va+len*PAGESIZ). Always succeeds.

void ae_unprotn(addr_t va, int len);

Make the region [va, va+len*PAGESIZ) writable. Fails if the current process does not have write access to any page.

void ae_tlbdeleten(addr_t va, int len);

Delete [va, va+len*PAGESIZ) from the TLB. Always succeeds.

void ae_tlbflush(void);

Flush all entries from STLB and TLB.

```
# Software refill of TLB: uses an STLB cache.

# 1. Compute hash function
mfc0    k0, c0_tlbcxt          # get virtual page  number and 11  bit process tag
mfc0    k1, c0_tlbcxt          # twice

#   Our hash function combines process tag with the lower bits of the virtual
#   page number (VPN) that missed:
#       (((c0_tlbcxt << 17 ^ c0_tlbctx) >> 16 ) & STLB_MASK&~7)
sll     k0, k0, 17                  # move VPN up
xor     k1, k0, k1                  # combine with process tag
srl     k1, k1, 16                  # move down (8 byte align)
andi    k0, k1, STLB_MASK & ~7  # remove upper and lower bits.

# 2. Index into STLB
lui     k1, HI(stlb)           # load STLB (at known location)
add     k1, k0, k1             # index into STLB

# 3. Load the physical page entry and STLB tag
lw      k0, LO(stlb)+4(k1)      # TLB entry
lw      k1, LO(stlb)+0(k1)      # STLB tag

# 4. Load TLB: we first load the fetched TLB entry into tlblo (but do not write
#    this register into the TLB); we then re  fetch tlbcxt in preparation to its
#    comparison to the STLB tag.
mtc0    k0, c0_tlblo           # (optimistically) load TLB entry
mfc0    k0, c0_tlbcxt          # get context again
nop                            # delay slot

# 5. Check tag (does not match   > jump to miss handler)
bne     k0, k1, stlb_miss       # compare tags to see if we got a hit
mfc0    k1, c0_epc             # get exception program counter

# 6. Tags matched: install entry into the TLB
tlbwr                       # write tlblo to TLB

# 7. Return from exception
j       k1                  # jump to resumption address
rfe                         # return from exception
```

Figure 9-7: Assembly code used by Aegis to lookup mapping in STLB (18 instructions).

```
/* LO portion of TLB mapping. */
struct lo {
        unsigned
                        w:1, /* write perm? */
                        :7,    /* reserved for software (libOS) */
                        g:1,  /* Global: TLB ignores the PID match req */
                        v:1, /* Valid: If not set, TLBL or TLBS miss occurs*/
                        d:1, /* Dirty */
                        n:1, /* Non  cacheable. */
                        pfn:20;  /* Page Frame Number: 31..12 of the pa */
};
```

Figure 9-8: Hardware defined “low” portion of a TLB entry (i.e., the part bound to a virtual page number).

Bibliography

[1] H. Abelson, G. J. Sussman, and J. Sussman. *Structure and Interpretation of Computer Programs*. MIT Press, 1996.

[2] M. Accetta, R. Baron, W. Bolosky, D. Golub, R. Rashid, A. Tevanian, and M. Young. Mach: a new kernel foundation for UNIX development. In *Proceedings of the Summer 1986 USENIX Conference*, pages 93–112, July 1986.

[3] T.E. Anderson. The case for application-specific operating systems. In *Third Workshop on Workstation Operating Systems*, pages 92–94, 1992.

[4] T.E. Anderson, B.N. Bershad, E.D. Lazowska, and H.M. Levy. Scheduler activations: Effective kernel support for the user-level management of parallelism. In *Proceedings of the Thirteenth ACM Symposium on Operating Systems Principles*, pages 95–109, October 1991.

[5] A.W. Appel and K. Li. Virtual memory primitives for user programs. In *Fourth International Conference on Architecture Support for Programming Languages and Operating Systems*, pages 96–107, Santa Clara, CA, April 1991.

[6] M. L. Bailey, B. Gopal, M. A. Pagels, L. L. Peterson, and P. Sarkar. PATHFINDER: A pattern-based packet classifier. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 115–123, Monterey, CA, USA, November 1994.

[7] K. Bala, M.F. Kaashoek, and W.E. Weihl. Software prefetching and caching for translation lookaside buffers. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 243–253, November 1994.

[8] J. Barrera. Invocation chaining: manipulating light-weight objects across heavy-weight boundaries. In *Proc. of 4th IEEE Workshop on Workstation Operating Systems*, pages 191–193, October 1993.

[9] B. N. Bershad, S. Savage, P. Pardyak, E. G. Sirer, M. Fiuczynski, D. Becker, S. Eggers, and C. Chambers. Extensibility, safety and performance in the SPIN operating system. In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, pages 267–284, Copper Mountain Resort, CO, USA, December 1995.

[10] B.N. Bershad, D.D. Redell, and J.R. Ellis. Fast mutual exclusion for uniprocessors. In *Proc. of the Conf. on Architectural Support for Programming Languages and Operating Systems*, pages 223–237, October 1992.

[11] Nathaniel Borenstein and James Gosling. Unix emacs: A retrospective. In *ACM SIGGRAPH Symposium on User Interface Software*, October 1988.

[12] E. Bugnion, S. Devine, and M. Rosenblum. Disco: running commodity operating systems on scalable multiprocessors. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, 1997.

[13] P. Cao, E. W. Felten, and K. Li. Implementation and performance of application-controlled file caching. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 165–178, November 1994.

[14] A. Chankhunthod, P. B. Danzig, C. Neerdaels, M. F. Sc hwartz, and K. J. Worrell. A hierarchical internet object cache. In *Proceedings of 1996 Usenix Technical Conference*, pages 153–163, January 1996.

[15] D. L. Chaum and R. S. Fabry. Implementing capability-based protection using encryption. Technical Report UCB/ERL M78/46, University of California at Berkeley, July 1978.

[16] D. Cheriton and K. Duda. A caching model of operating system kernel functionality. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 179–193, November 1994.

[17] D. R. Cheriton. An experiment using registers for fast message-based interprocess communication. *Operating Systems Review*, 18:12–20, October 1984.

[18] D. D. Clark and D. L. Tennenhouse. Architectural considerations for a new generation of protocols. In *ACM Communication Architectures, Protocols, and Applications (SIGCOMM) 1990*, pages 200–208, Philadelphia, PA, USA, September 1990.

[19] D.D. Clark. On the structuring of systems using upcalls. In *Proceedings of the Tenth ACM Symposium on Operating Systems Principles*, pages 171–180, December 1985.

[20] R. J. Creasy. The origin of the VM/370 time-sharing system. *IBM J. Research and Development*, 25(5):483–490, September 1981.

[21] H. Custer. *Inside Windows/NT*. Microsoft Press, Redmond, WA, 1993.

[22] P. Deutsch and C. A. Grant. A flexible measurement tool for software systems. *Information Processing 71*, 1971.

[23] P. Druschel, L. L. Peterson, and B. S. Davie. Experiences with a high-speed network adaptor: A software perspective. In *ACM Communication Architectures, Protocols, and Applications (SIGCOMM) 1994*, pages 2–13, London, UK, August 1994.

[24] A. Edwards, G. Watson, J. Lumley, D. Banks, C. Clamvokis, and C. Dalton. User-space protocols deliver high performance to applications on a low-cost Gb/s LAN. In *ACM Communication Architectures, Protocols, and Applications (SIGCOMM) 1994*, pages 14–24, London, UK, August 1994.

[25] D. R. Engler, M. F. Kaashoek, and J. O' Toole Jr. Exokernel: an operating system architecture for application-specific resource management. In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, pages 251–266, Copper Mountain Resort, Colorado, December 1995.

[26] D. R. Engler, M. F. Kaashoek, and J. O' Toole. The operating system kernel as a secure programmable machine. In *Proceedings of the Sixth SIGOPS European Workshop*, pages 62–67, September 1994.

[27] D. R. Engler, M. F. Kaashoek, and J. O' Toole. The operating system kernel as a secure programmable machine. In *Operating systems review*, January 1995.

[28] D. R. Engler, D. Wallach, and M. F. Kaashoek. Efficient, safe, application-specific message processing. Technical Memorandum MIT/LCS/TM533, MIT, March 1995.

[29] Dawson R. Engler. Simple, robust online verification of program correctness. Submitted for publication.

[30] Dawson R. Engler. Efficient verification of demonically-implemented integer functions (or, demonic determinism, trusted results). available on request, December 1997.

[31] D.R. Engler and M.F. Kaashoek. DPF: fast, flexible message demultiplexing using dynamic code generation. In *ACM Communication Architectures, Protocols, and Applications (SIGCOMM) 1996*, pages 53–59, Stanford, CA, USA, August 1996.

[32] Marc Fiuczynski and Brian Bershad. An extensible protocol architecture for application-specific networking. In *Proceedings of the 1996 Winter USENIX Conference*, pages 55–64, January 1996.

[33] B. Ford, K. Van Maren, J. Lepreau, S. Clawson, B. Robinson, and Jeff Turner. The FLUX OS toolkit: Reusable components for OS implementation. In *Proc. of Sixth Workshop on Hot Topics in Operating Systems*, pages 14–19, May 1997.

[34] Bryan Ford, Mike Hibler, Jay Lepreau, Patrick Tullman, Godmar Back, and Steven Clawson. Microkernels meet recursive virtual machines. In *Proceedings of the Second Symposium on Operating System Design and Implementation (OSDI 1996)*, October 1996.

[35] Bryan Ford and Sai R. Susarla. CPU inheritance scheduling. In *Proceedings of the Second Symposium on Operating System Design and Implementation (OSDI 1996)*, October 1996.

[36] G. Ganger and Y. Patt. Metadata update performance in file systems. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 49–60, November 1994.

[37] Gregory R. Ganger and M. Frans Kaashoek. Embedded inodes and explicit grouping: Exploiting disk bandwidth for small files. In *Proceedings of the 1997 USENIX Technical Conference*, 1997.

[38] R. P. Goldberg. Survey of virtual machine research. *IEEE Computer*, pages 34–45, June 1974.

[39] D. Golub, R. Dean, A. Forin, and R. Rashid. UNIX as an application program. In *USENIX 1990 Summer Conference*, pages 87–95, June 1990.

[40] J. Gosling. Java intermediate bytecodes. In *Proc. of ACM SIGPLAN workshop on Intermediate Representations*, pages 111–118, march 1995.

[41] P. Brinch Hansen. The nucleus of a multiprogramming system. *Communications of the ACM*, 13(4):238–241, April 1970.

[42] H. Härtig, M. Hohmuth, J. Liedtke, and S. Schönberg andJ. Wolter. The performance of μ-kernel-based systems. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, 1997.

[43] J.H. Hartman, A.B. Montz, D. Mosberger, S.W. O'Malley, L.L. Peterson, and T.A. Proebsting. Scout: A communication-oriented operating system. Technical Report TR 94-20, University of Arizona, Tucson, AZ, June 1994.

[44] K. Harty and D.R. Cheriton. Application-controlled physical memory using external page-cache management. In *Fifth International Conference on Architecture Support for Programming Languages and Operating Systems*, pages 187–199, October 1992.

[45] D. Hitz. An NFS file server appliance. Technical Report 3001, Network Applicance Corporation, March 1995.

[46] J. Huck and J. Hays. Architectural support for translation table management in large address space machines. In *Proceedings of the 19th International Symposium on Computer Architecture*, pages 39–51, May 1992.

[47] Norman P. Jouppi. Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers. In *17th Annual International Symposium on Computer Architecture*, pages 364–373, May 1990.

[48] M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Hector M. Briceno, Russell Hunt, David Mazieres, Thomas Pinckney, Robert Grimm, John Jannotti, and Kenneth Mackenzie. Application performance and flexibility on exokernel systems. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, October 1997.

[49] M.F. Kaashoek, D.R. Engler, D.H. Wallach, and G. Ganger. Server operating systems. In *SIGOPS European Workshop*, pages 141–148, September 1996.

[50] Gerry Kane and Joe Heinrich. *MIPS RISC Architecture*. Prentice Hall, 1992.

[51] K. Krueger, D. Loftesness, A. Vahdat, and T. Anderson. Tools for development of application-specific virtual memory management. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA) 1993*, pages 48–64, October 1993.

[52] B. W. Lampson. Hints for computer system design. In *Proceedings of the Eighth ACM Symposium on Operating Systems Principles*, pages 33–48, December 1983.

[53] B.W. Lampson. On reliable and extendable operating systems. *State of the Art Report, Infotech*, 1, 1971.

[54] B.W. Lampson and R.F. Sproull. An open operating system for a single-user machine. *Proceedings of the Seventh ACM Symposium on Operating Systems Principles*, pages 98–105, December 1979.

[55] C.H. Lee, M.C. Chen, and R.C. Chang. HiPEC: high performance external virtual memory caching. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 153–164, 1994.

[56] Ian Leslie, Derek McAuley, Richard Black, Timothy Roscoe, Paul Barham, David Evers, Robin Fairbairns, , and Eoin Hyden. The design and implementation of an operating system to support distributed multimedia applications. *IEEE Journal on selected areas in communication*, 14(7):1280–1297, September 1996.

[57] J. Liedtke. Improving IPC by kernel design. In *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, pages 175–188, December 1993.

[58] J. Liedtke. On micro-kernel construction. In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, December 1995.

[59] Steve Lucco. Personal communication. Use of undocumented proprietary formats as a technique to impede third-party additions, August 1997.

[60] Kenneth Mackenzie, John Kubiatowicz, Matthew Frank, Walter Lee, Victor Lee, Anant Agarwal, and M. Frans Kaashoek. UDM: User Direct Messaging for General-Purpose Multiprocessing. Technical Memo MIT/LCS/TM-556, March 1996.

[61] C. Maeda and B. N. Bershad. Protocol service decomposition for high-performance networking. In *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, pages 244–255, Asheville, NC, USA, 1993.

[62] David Mazieres and M. Frans Kaashoek. Secure applications need flexibile operating systems. In *HotOS-VI*, 1997.

[63] David Mazieres and M. Frans Kaashoek. Secure applications need flexible operating systems. In *Proceedings of the6th Workshop on Hot Topics in Operating Systems*, May 1997.

[64] S. McCanne and V. Jacobson. The BSD packet filter: A new architecture for user-level packet capture. In *USENIX Technical Conference Proceedings*, pages 259–269, San Diego, CA, Winter 1993. USENIX.

[65] J.C. Mogul, R.F. Rashid, and M.J. Accetta. The packet filter: An efficient mechanism for user-level network code. In *Proceedings of the Eleventh ACM Symposium on Operating Systems Principles*, pages 39–51, Austin, TX, USA, November 1987.

[66] A. C. Myers and B. Liskov. Decentralized model for information flow control. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, October 1997.

[67] D. Nagle, R. Uhlig, T. Stanley, S. Sechrest, T. Mudge, and R. Brown. Design tradeoffs for software-managed TLBs. In *20th Annual International Symposium on Computer Architecture*, pages 27–38, May 1993.

[68] NCSA, University of Illinois, Urbana-Champaign. NCSA HTTPd. http://hoohoo.ncsa.uiuc.edu/index.html.

[69] J. K. Ousterhout. Why aren't operating systems getting faster as fast as hardware? In *Proceedings of the Summer 1990 USENIX Conference*, pages 247–256, June 1990.

[70] V. Pai, P. Druschel, and W. Zwaenepoel. I/O-lite: a unified I/O buffering and caching system. Technical Report *http://www.cs.rice.edu/ vivek/IO-lite.html*, Rice University, 1997.

[71] Przemyslaw Pardyak and Brian Bershad. Dynamic binding for an extensible system. In *Proceedings of the Second USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 201–212, October 1996.

[72] R. Hugo Patterson, Garth A. Gibson, Eka Ginting, Daniel Stodolsky, and Jim Zelenka. Informed prefetching and caching. In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, Copper Mountain Resort, CO, December 1995.

[73] D. Probert, J.L. Bruno, and M. Karzaorman. SPACE: A new approach to operating system abstraction. In *International Workshop on Object Orientation in Operating Systems*, pages 133–137, October 1991.

[74] J.S. Quarterman, A. Silberschatz, and J.L. Peterson. 4.2BSD and 4.3BSD as examples of the UNIX system. *Computing Surveys*, 17(4):379–418, December 1985.

[75] D.D. Redell, Y.K. Dalal, T.R. Horsley, H.C. Lauer, W.C. Lynch, P.R. McJones, H.G. Murray, and S.C. Purcell. Pilot: An operating system for a personal computer. *Communications of the ACM*, 23(2):81–92, February 1980.

[76] Timothy Roscoe. *The Structure of a Multi-Service Operating System*. Phd Thesis, Technical Report 376, Cambridge, 1995.

[77] M. Rozier, V. Abrossimov, F. Armand, I. Boule, M. Gien, M. Guillemont, F. Herrmann, C. Kaiser, S. Langlois, P. Leonard, and W. Neuhauser. Chorus distributed operating system. *Computing Systems*, 1(4):305–370, 1988.

[78] M. Seltzer, Y. Endo, C. Small, and K. Smith. Dealing with disaster: Surviving misbehaved kernel extensions. In *Proceedings of the Second Symposium on Operating Systems Design and Implementation*, pages 213–228, October 1996.

[79] C. Small and M. Seltzer. Vino: an integrated platform for operating systems and database research. Technical Report TR-30-94, Harvard, 1994.

[80] Christopher Small and Margo Seltzer. A comparison of os extension technologies. In *Proceedings of the 1996 USENIX Conference*, 1996.

[81] Richard Stallman. Emacs, the extensible, customizable self-documenting display editor. In *ACM SIGPLAN SIGOA Symposium on Text Manipulation*, June 1981.

[82] V. Buch T. von Eicken, A. Basu and W. Vogels. U-Net: A user-level network interface for parallel and distributed computing. In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, pages 40–53, Copper Mountain Resort, CO, USA, 1995.

[83] Madhusudhan Talluri, Mark D. Hill, and Yousef A. Khalidi. A new page table for 64-bit address spaces. In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, Copper Mountain Resort, Colorado, December 1995.

[84] A.S. Tanenbaum, R. van Renesse, H. van Staveren, G. Sharp, S.J. Mullender, A. Jansen, and G. van Rossum. Experiences with the Amoeba distributed operating system. *Communications of the ACM*, 33(12):46–63, December 1990.

[85] D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. Til: A type-directed optimizing compiler for ml. In *Proceedings of the SIGPLAN '96 Conference on Programming Language Design and Implementation*, Philadelphia, PA, May 1996.

[86] D.L. Tennenhouse and David J. Wetherall. Towards an active network architecture. In *Proc. Multimedia, Computing, and Networking 96*, January 1996.

[87] C. A. Thekkath and H. M. Levy. Hardware and software support for efficient exception handling. In *Sixth International Conference on Architecture Support for Programming Languages and Operating Systems*, pages 110–121, October 1994.

[88] C. A. Thekkath, H. M. Levy, and E. D. Lazowska. Separating data and control transfer in distributed operating systems. In *Sixth International Conference on Architecture Support for Programming Languages and Operating Systems*, pages 2–11, San Francisco, California, October 1994.

[89] R. Wahbe, S. Lucco, T. Anderson, and S. Graham. Efficient software-based fault isolation. In *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, pages 203–216, Asheville, NC, USA, December 1993.

[90] C. A. Waldspurger and W. E. Weihl. Lottery scheduling: Flexible proportional-share resource management. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 1–11, November 1994.

[91] C. A. Waldspurger and W. E. Weihl. Stride scheduling: deterministic proportional-share resource management. Technical Memorandum MIT/LCS/TM528, MIT, June 1995.

[92] D. A. Wallach, D. R. Engler, and M. F. Kaashoek. ASHs: Application-specific handlers for high-performance messaging. In *ACM Communication Architectures, Protocols, and Applications (SIGCOMM '96)*, Stanford, California, August 1996.

[93] Deborah A. Wallach. *Supporting application-specific libraries for communication*. PhD thesis, M.I.T., 1996.

[94] W. Wulf, E. Cohen, W. Corwin, A. Jones, R. Levin, C. Pierson, and F. Pollack. HYDRA: The kernel of a multiprocessing operating system. *Communications of the ACM*, 17(6):337–345, July 1974.

[95] M. Yuhara, B. Bershad, C. Maeda, and E. Moss. Efficient packet demultiplexing for multiple endpoints and large messages. In *Proceedings of the Winter 1994 USENIX Conference*, pages 153–165, San Francisco, CA, USA, January 1994.